arXiv:1407.7759v1 [cs.DS] 29 Jul 2014

Simultaneous Approximation of Constraint Satisfaction Problems

Amey Bhangale * Swastik Kopparty † Sushant Sachdeva ‡

July 30, 2014

Abstract

Given k collections of 2SAT clauses on the same set of variables V, can we find one assignment that satisfies a large fraction of clauses from *each* collection? We consider such *simultaneous* constraint satisfaction problems, and design the first nontrivial approximation algorithms in this context.

Our main result is that for every CSP $\mathcal{F}$, for $k < \tilde{O}(\log^{1/4} n)$, there is a polynomial time constant factor *Pareto* approximation algorithm for k simultaneous Max-$\mathcal{F}$-CSP instances. Our methods are quite general, and we also use them to give an improved approximation factor for simultaneous Max-w-SAT (for $k < \tilde{O}(\log^{1/3} n)$). In contrast, for $k = \omega(\log n)$, no nonzero approximation factor for k simultaneous Max-$\mathcal{F}$-CSP instances can be achieved in polynomial time (assuming the Exponential Time Hypothesis).

These problems are a natural meeting point for the theory of constraint satisfaction problems and multiobjective optimization. We also suggest a number of interesting directions for future research.

*Department of Computer Science. Rutgers University. Research supported in part by NSF grant CCF-1253886. amey.bhangale@rutgers.edu

†Department of Mathematics & Department of Computer Science. Rutgers University. Research supported in part by a Sloan Fellowship and NSF grant CCF-1253886. swastik.kopparty@rutgers.edu

‡Department of Computer Science, Yale University. Research supported by the NSF grants CCF-0832797, CCF-1117309, and Daniel Spielman's & Sanjeev Arora's Simons Investigator Grants. Part of this work was done when this author was at the Simons Institute for the Theory of Computing, UC Berkeley, and at the Department of Computer Science, Princeton University. Email: sachdeva@cs.yale.edu

1 Introduction

The theory of approximation algorithms for constraint satisfaction problems (CSPs) is a very central and well developed part of modern theoretical computer science. Its study has involved fundamental theorems, ideas, and problems such as the PCP theorem, linear and semidefinite programming, randomized rounding, the Unique Games Conjecture, and deep connections between them [AS98, ALM+98, GW95, Kho02, Rag08, RS09].

In this paper, we initiate the study of *simultaneous approximation algorithms* for constraint satisfaction problems. A typical such problem is the simultaneous MAX-CUT problem: Given a collection of k graphs $G_i = (V, E_i)$ on the same vertex set V, the problem is to find a single cut (i.e., a partition of V) so that in *every* G_i, a large fraction of the edges go across the cut.

More generally, let q be a constant positive integer, and let $\mathcal{F}$ be a set of bounded-arity predicates on $[q]$-valued variables. Let V be a set of n $[q]$-valued variables. An $\mathcal{F}$-CSP is a weighted collection $\mathcal{W}$ of constraints on V, where each constraint is an application of a predicate from $\mathcal{F}$ to some variables from V. For an assignment $f : V \to [q]$ and a $\mathcal{F}$-CSP instance $\mathcal{W}$, we let $\mathsf{val}(f, \mathcal{W})$ denote the total weight of the constraints from $\mathcal{W}$ satisfied by f. The MAX-$\mathcal{F}$-CSP problem is to find f which maximizes $\mathsf{val}(f, \mathcal{W})$. If $\mathcal{F}$ is the set of all predicates on $[q]$ of arity w, then MAX-$\mathcal{F}$-CSP is also called MAX-w-CSP$_q$.

We now describe the setting for the problem we consider: *k-fold simultaneous* MAX-$\mathcal{F}$-CSP. Let $\mathcal{W}_1, \ldots, \mathcal{W}_k$ be $\mathcal{F}$-CSPs on V, each with total weight 1. Our high level goal is to find an assignment $f : V \to [q]$ for which $\mathsf{val}(f, \mathcal{W}_\ell)$ is large for all $\ell \in [k]$.

These problems fall naturally into the domain of multi-objective optimization: there is a common search space, and multiple objective functions on that space. Since even optimizing one of these objective functions could be NP-hard, it is natural to resort to approximation algorithms. Below, we formulate some of the approximation criteria that we will consider, in decreasing order of difficulty:

1. **Pareto approximation:** Suppose $(c_1, \ldots, c_k) \in [0, 1]^k$ is such that there is an assignment f^* with $\mathsf{val}(f^*, \mathcal{W}_\ell) \geqslant c_\ell$ for each $\ell \in [k]$.

   An α-Pareto approximation algorithm in this context is an algorithm, which when given $(c_1, \ldots, c_k)$ as input, finds an assignment f such that $\mathsf{val}(f, \mathcal{W}_\ell) \geqslant \alpha \cdot c_\ell$, for each $\ell \in [k]$.

2. **Minimum approximation:** This is basically the Pareto approximation problem when $c_1 = c_2 = \ldots = c_k$. Define OPT to be the maximum, over all assignments f^*, of $\min_{\ell \in [k]} \mathsf{val}(f^*, \mathcal{W}_\ell)$.

   An α-minimum approximation algorithm in this context is an algorithm which finds an assignment f such that $\min_{\ell \in [k]} \mathsf{val}(f, \mathcal{W}_\ell) \geqslant \alpha \cdot \text{OPT}$.

3. **Detecting Positivity:** This is a very special case of the above, where the goal is simply to determine whether there is an assignment f which makes $\mathsf{val}(f, \mathcal{W}_\ell) > 0$ for all $\ell \in [k]$.

   At the surface, this problem appears to be a significant weakening of the the simultaneous approximation goal.

When $k = 1$, minimum approximation and Pareto approximation correspond to the classical MAX-CSP approximation problems (which have received much attention). Our focus in this paper is on general k. As we will see in the discussions below, the nature of the problem changes quite a bit for $k > 1$. In particular, direct applications of classical techniques like random assignments and convex programming relaxations fail to give even a constant factor approximation.

The theory of *exact* multiobjective optimization has been very well studied, (see eg. [PY00, Dia11] and the references therein). For several optimization problems such as shortest paths, minimum spanning trees, matchings, etc, there are polynomial time algorithms that solve the multiobjective versions exactly. For Max-SAT, simultaneous approximation was studied by Glaßer *et al.* [GRW11].

We have two main motivations for studying simultaneous approximations for CSPs. Most importantly, these are very natural algorithmic questions, and capture naturally arising constraints in a way which more naïve formulations (such as taking linear combinations of the given CSPs) cannot. Secondly, the study of simultaneous approximation algorithms for CSPs sheds new light on various aspects of standard approximation algorithms for CSPs. For example, our algorithms are able to favorably exploit some features of the trivial random-assignment-based 1/2-approximation algorithm for Max-CUT, that are absent in the more sophisticated SDP-based 0.878-approximation algorithm of Goemans-Williamson [GW95].

1.1 Observations about simultaneous approximation

We now discuss why a direct application of the classical CSP algorithms fails in this setting, and limitations on the approximation ratios that can be achieved.

We begin with a trivial remark. Finding an α-minimum (or Pareto) approximation to the k-fold Max-$\mathcal{F}$-CSP is at least as hard as finding an α-approximation the classical Max-$\mathcal{F}$-CSP problem (i.e., $k = 1$). Thus the known limits on polynomial-time approximability extend naturally to our setting.

Max-1-SAT. The simplest simultaneous CSP is Max-1-SAT. The problem of getting a 1-Pareto or 1-minimum approximation to k-fold simultaneous Max-1-SAT is essentially the **NP**-hard SUBSET-SUM problem. There is a simple $2^{\mathrm{poly}(k/\varepsilon)} \cdot \mathrm{poly}(n)$-time $(1-\varepsilon)$-Pareto approximation algorithm based on dynamic programming.

It is easy to see that detecting positivity of a k-fold simultaneous Max-1-SAT is exactly the same problem as detecting satisfiability of a SAT formula with k clauses (a problem studied in the fixed parameter tractability community. Thus, this problem can be solved in time $2^{O(k)} \cdot \mathrm{poly}(n)$ (see [Mar13]), and under the Exponential Time Hypothesis, one does not expect a polynomial time algorithm when $k = \omega(\log n)$.

Random Assignments. Let us consider algorithms based on random assignments. A typical example is Max-CUT. A uniformly random cut in a weighted graph graph cuts 1/2 the total weight in expectation. This gives a 1/2-approximation to the classical Max-CUT problem.

If the cut value is concentrated around 1/2, with high probability, we would obtain a cut that's simultaneously good for all instances. For an *unweighted* graph[1] G with $\omega(1)$ edges, a simple variance calculation shows that a uniformly random cut in the graph cuts a $\left(\frac{1}{2} - o(1)\right)$ fraction of the edges *with high probability*. Thus by a union bound, for $k = O(1)$ simultaneous unweighted instances $G_1, \ldots, G_k$ of Max-CUT, a uniformly random cut gives a $\left(\frac{1}{2} - o(1)\right)$-minimum (and Pareto) approximation with high probability. However, for *weighted* graphs, the concentration no longer holds, and the algorithm fails to give any constant factor approximation.

[1] We use the term "unweighted" to refer to instances where all the constraints have the same weight. When we talk about simultaneous approximation for unweighted instances $\mathcal{W}_1, \ldots, \mathcal{W}_k$ of MAX-$\mathcal{F}$-CSP, we mean that in each instance $\mathcal{W}_i$, all constraints with nonzero weight have the equal weights (but that equal weight can be different for different i).

For general CSPs, even for unweighted instances, the total weight satisfied by a random assignment does not necessarily concentrate. In particular, there is no "trivial" random-assignment-based constant factor approximation algorithm for simultaneous general CSPs.

SDP Algorithms. How do algorithms based on semi-definite programming (SDP) generalize to the simultaneous setting?

For the usual MAX-CUT problem ($k = 1$), the celebrated Goemans-Williamson SDP algorithm [GW95] gives a 0.878-approximation. The SDP relaxation generalizes naturally to to the simultaneous setting; it allows us to find a vector solution which is a simultaneously good cut for $G_1, \ldots, G_k$. Perhaps we apply hyperplane rounding to the SDP solution to obtain a simultaneously good cut for all G_i? We know that each G_i gets a good cut in expectation, but we need each G_i to get a good cut *with high probability* to guarantee a simultaneously good cut.

However, there are cases where the hyperplane rounding fails completely. For weighted instances, the SDP does not have *any* constant integrality gap. For unweighted instances, for every fixed k, we find an instance of k-fold simultaneous MAX-CUT (with arbitrarily many vertices and edges) where the SDP relaxation has value $1 - \Omega\left(\frac{1}{k^2}\right)$, while the optimal simultaneous cut has value only $1/2$. Furthermore, applying the hyperplane rounding algorithm to this vector solution gives (with probability 1) a simultaneous cut value of *0*. These integrality gaps are described in Section C.

Thus the natural extension of SDP based techniques for simultaneous approximation fail quite spectacularly. A-priori, this failure is quite surprising, since SDPs (and LPs) generalize to the multiobjective setting seamlessly.

Matching Random Assignments? Given the ease and simplicity of algorithms based on random assignments for $k = 1$, giving algorithms in the simultaneous setting that match their approximation guarantees is a natural benchmark. Perhaps it is always possible to do as well in the simultaneous setting as a random assignment for one instance?

Somewhat surprisingly, this is incorrect. For simultaneous MAX-Ew-SAT (CNF-SAT where every clause has exactly w distinct literals), a simple reduction from MAX-E3-SAT (with $k = 1$) shows that it is **NP**-hard to give a $(7/8 + \varepsilon)$-minimum approximation for k-fold simultaneous MAX-Ew-SAT for large enough constants k.

Proposition 1.1 *For all integers $w \geqslant 4$ and $\varepsilon > 0$, given $k \geqslant 2^{w-3}$ instances of* MAX-Ew-SAT *that are simultaneously satisfiable, it is* **NP***-hard to find a $(7/8 + \varepsilon)$-minimum (or Pareto) approximation.*

On the other hand, a random assignment to a single MAX-Ew-SAT instance satisfies a $1 - 2^{-w}$ fraction of constraints in expectation.

This shows that simultaneous CSPs can have worse approximation factors than that expected from a random assignment. In particular, it shows that simultaneous CSPs can have worse approximation factors than their classical ($k = 1$) counterparts.

1.2 Results

Our results address the approximability of k-fold simultaneous MAX-$\mathcal{F}$-CSP for large k. Our main algorithmic result shows that for every $\mathcal{F}$, and k not too large, k-fold simultaneous MAX-$\mathcal{F}$-CSP has a constant factor Pareto approximation algorithm.

Theorem 1.2 *Let q, w be constants. Then for every $\varepsilon > 0$, there is a $2^{O(k^4/\varepsilon^2 \log(k/\varepsilon))} \cdot \mathrm{poly}(n)$-time $\left(\frac{1}{q^{w-1}} - \varepsilon\right)$-Pareto approximation algorithm for k-fold simultaneous* Max-w-CSP$_q$.

The dependence on k implies that the algorithm runs in polynomial time up to $k = \tilde{O}((\log n)^{1/4})$ simultaneous instances [2]. The proof of the above Theorem appears in Section 4, and involves a number of ideas. In order to make the ideas clearer, we first describe the main ideas for approximating simultaneous Max-2-AND (which easily implies the $q = w = 2$ special case of the above theorem); this appears in Section 3.

For particular CSPs, our methods allow us to do significantly better, as demonstrated by our following result for Max-w-SAT.

Theorem 1.3 *Let w be a constant. For every $\varepsilon > 0$, there is a $2^{O(k^3/\varepsilon^2 \log(k/\varepsilon))} \cdot \mathrm{poly}(n)$-time $(3/4 - \varepsilon)$-Pareto approximation algorithm for k-fold* Max-w-SAT.

Given a single Max-Ew-SAT instance, a random assignment satisfies a $1 - 2^{-w}$ fraction of the constraints in expectation. The approximation ratio achieved by the above theorem seems unimpressive in comparison (even though it is for general Max-w-SAT). However, Proposition 1.1 demonstrates it is **NP**-hard to do much better.

Remarks

1. As demonstrated by Proposition 1.1, it is sometimes impossible to match the approximation ratio achieved by a random assignment for $k = 1$. By comparison, the approximation ratio given by Theorem 1.2 is slightly better than that achieved by a random assignment ($1/q^w$). This is comparable to the best possible approximation ratio for $k = 1$, which is w/q^{w-1} up to constants [MM12, Cha13]. Our methods also prove that picking the best assignment out of $2^{O(k^4/\varepsilon^2 \log(k/\varepsilon))}$ independent and uniformly random assignments achieves a $(1/q^w - \varepsilon)$-Pareto approximation with high probability.

2. Our method is quite general. For any CSP with a convex relaxation and an associated rounding algorithm that assigns each variable independently from a distribution with certain smoothness properties (see Section 3.2), it can be combined with our techniques to achieve essentially the same approximation ratio for k simultaneous instances.

3. We reiterate that Pareto approximation algorithms achieve a multiplicative approximation for each instance. One could also consider the problem of achieving simultaneous approximations with an α-multiplicative and ε-additive error. This problem can be solved by a significantly simpler algorithm and analysis (but note that this variation does not even imply an algorithm for detecting positivity).

1.3 Complementary results

1.3.1 Refined hardness results

As we saw earlier, assuming ETH, there is no algorithm for even detecting positivity of k-fold simultaneous Max-1-SAT for $k = \omega(\log n)$. There are trivial examples of CSPs for which detecting

[2] The $\tilde{O}(\cdot)$ hides $\mathrm{poly}(\log \log n)$ factors.

positivity (and in fact 1-Pareto approximation) can be solved efficiently: eg. simultaneous CSPs based on monotone predicates (where no negations of variables are allowed) are maximally satisfied by the all-1s assignment. Here we prove that for any "nontrivial" collection of Boolean predicates $\mathcal{F}$, assuming ETH, there is no polynomial time algorithm for detecting positivity for k-fold simultaneous MAX-$\mathcal{F}$-CSP instances for $k = \omega(\log n)$. In particular, it is hard to obtain any poly-time constant factor approximation for $k = \omega(\log n)$. This implies a complete *dichotomy theorem* for constant factor approximations of k-fold simultaneous Boolean CSPs.

A predicate $P : \{0,1\}^w \to \{\text{TRUE}, \text{FALSE}\}$ is said to be 0-*valid*/1-*valid* if the all-0-assignment/all-1-assignment satisfies P. We call a collection $\mathcal{F}$ of predicates 0-*valid*/1-*valid* if all predicates in $\mathcal{F}$ are 0-valid/1-valid. Clearly, if $\mathcal{F}$ is 0-valid or 1-valid, the simultaneous MAX-$\mathcal{F}$-CSP instances can be solved exactly (by considering the all-0-assignment/all-1-assignment). Our next theorem shows that detecting positivity of $\omega(\log n)$-fold simultaneous MAX-$\mathcal{F}$-CSP , for all other $\mathcal{F}$, is hard.

Theorem 1.4 *Assume the Exponential Time Hypothesis [IP01, IPZ01]. Let $\mathcal{F}$ be a fixed finite set of Boolean predicates. If $\mathcal{F}$ is not* 0-*valid or* 1-*valid, then for $k = \omega(\log n)$, detecting positivity of k-fold simultaneous* MAX-$\mathcal{F}$-CSP *on n variables requires time super-polynomial in n.*

Crucially, this hardness result holds even if we require that every predicate in an instance has all its inputs being *distinct variables*.

Our proof uses techniques underlying the dichotomy theorems of Schaefer [Sch78] for exact CSPs, and of Khanna *et al.* [KSTW01] for MAX-CSPs (although our easiness criterion is different from the easiness criteria in both these papers).

1.3.2 Simultaneous approximations via SDPs

It is a tantalizing possibility that one could use SDPs to improve the LP-based approximation algorithms that we develop. Especially for constant k, it is not unreasonable to expect that one could obtain a constant factor Pareto or minimum approximation, for k-fold simultaneous CSPs, better than what can be achieved by linear programming methods.

In this direction, we show how to use simultaneous SDP relaxations to obtain a polynomial time $(1/2 + \Omega(1/k^2))$-minimum approximation for k-fold simultaneous MAX-CUT on *unweighted graphs*.

Theorem 1.5 *For large enough n, there is an algorithm that, given k-fold simultaneous unweighted* MAX-CUT *instances on n vertices, runs in time $2^{2^{2^{O(k)}}} \cdot \text{poly}(n)$, and computes a $\left(\frac{1}{2} + \Omega\left(\frac{1}{k^2}\right)\right)$-minimum approximation.*

1.4 Our techniques

For the initial part of this discussion, we focus on the $q = w = 2$ case, and only achieve a $1/4 - \varepsilon$ Pareto approximation.

Preliminary Observations First let us analyze the behavior of the uniformly random assignment algorithm. It is easy to compute, for each instance $\ell \in [k]$, the expected weight of satisfied constraints in instance ℓ, which will be at least $\frac{1}{4}$ of the total weight all constraints in instance ℓ. If we knew for some reason that in each instance the weight of satisfied constraints was concentrated around this expected value *with high probability*, then we could take a union bound over all the instances and conclude that a random assignment satisfies many constraints in each instance with

high probability. It turns out that for any instance where the desired concentration does not occur, there is some variable in that instance which has high degree (i.e., the weight of all constraints involving that variable is a constant fraction of the total weight of all constraints). Knowing that there is such a high degree variable seems very useful for our goal of finding a good assignment, since we can potentially influence the satisfaction of the instance quite a bit by just by changing this one variable.

This motivates a high-level plan: either proceed by using the absence of influential variables to argue that a random assignment will succeed, or proceed by trying to set the influential variables.

An attempt The above high-level plan motivates the following high-level algorithm. First we identify a set $S \subseteq V$ of "influential" variables. This set of influential variables should be of small ($O(\log n)$) size, so that we can try out all assignments to these variables. Next, we take a random assignment to the remaining variables, $g : V \backslash S \to \{0,1\}$. Finally, for each possible assignment $h : S \to \{0,1\}$, we consider the assignment $h \cup g : V \to \{0,1\}$ as a candidate solution for our simultaneous CSP. We output the assignment, if any, that has $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant \alpha \cdot c_\ell$ for each $\ell \in [k]$. This concludes the description of the high-level algorithm.

For the analysis, we would start with the ideal assignment $f^* : V \to \{0,1\}$ achieving $\mathsf{val}(f^*, \mathcal{W}_\ell) \geqslant c_\ell$ for each $\ell \in [k]$. Consider the step of the algorithm where h is taken to equal $h^* \stackrel{\text{def}}{=} f^*|_S$. We would like to say that for each $\ell \in [k]$ we have:

$$\mathsf{val}(h^* \cup g, \mathcal{W}_\ell) \geqslant (\frac{1}{4} - \varepsilon) \cdot \mathsf{val}(f^*, \mathcal{W}_\ell),$$

with high probability, when $g : V \backslash S \to \{0,1\}$ is chosen uniformly at random. (We could then conclude the analysis by a union bound.)

A simple calculation shows that $\mathbf{E}[\mathsf{val}(h^* \cup g, \mathcal{W}_\ell)] \geqslant \frac{1}{4} \cdot \mathsf{val}(f^*, \mathcal{W}_\ell)$, so each instance is well satisfied in expectation. Our hope is thus that $\mathsf{val}(h^* \cup g, \mathcal{W}_\ell)$ is concentrated around its mean with high probability.

There are two basic issues with this approach[3]:

1. The first issue is how to define the set S of influential variables. For some special CSPs (such as Max-Cut and Max-Sat), there is a natural choice which works (to choose a set of variables with high degree, which is automatically small). But for general CSPs, it could be the case that variables with exponentially small degree are important contributors to the ideal assignment f^*.

2. Even if one chooses the set S of influential variables appropriately, the analysis cannot hope to argue that $\mathsf{val}(h^* \cup g, \mathcal{W}_\ell)$ concentrates around its expectation with high probability. Indeed, it can be the case that for a random assignment g, $\mathsf{val}(h^* \cup g, \mathcal{W}_\ell)$ is not concentrated at all.

A working algorithm: Our actual algorithm and analysis solve these problems by proceeding in a slightly different way. The first key idea is to find the set of influential variables by iteratively

[3] These problems do not arise if we only aim for the weaker "additive-multiplicative" Pareto approximation guarantee (where one allows for both some additive loss and multiplicative loss in the approximation), and in fact the above mentioned high-level plan does work. The pure multiplicative approximation guarantee seems to be significantly more delicate.

including variables into this set, and simultaneously assigning these variables. This leads to a tree-like evolution of the set of influential variables. The second key idea is in the analysis: instead of arguing about the performance of the algorithm when considering the partial assignment $h^* = f^*|_S$, we will perform a delicate *perturbation* of h^* to obtain an $h' : S \to \{0,1\}$, and show that $\mathsf{val}(h' \cup g, \mathcal{W}_\ell)$ is as large as desired. Intuitively, this perturbation only slightly worsens the satisfied weight of h^*, while reducing the reliance of the good assignment f^* on any specialized properties of $f^*|_S$.

To implement this, the algorithm will maintain a tree of possible evolutions of a set $S \subseteq V$ and a partial assignment $\rho : S \to \{0,1\}$. In addition, every variable $x \in S$ will be labelled by an instance $\ell \in [k]$. The first stage of the algorithm will grow this tree in several steps. In the beginning, at the root of the tree, we have $S = \varnothing$. At every stage, we will either terminate that branch of the tree, or else increase the size of the set S by 1 (or 2), and consider all 2 (or 4) extensions of ρ to the newly grown S.

To grow the tree, the algorithm considers a random assignment $g : V \backslash S \to \{0,1\}$, and computes, for each instance $i \in [k]$, the expected satisfied weight $\mathbf{E}_g[\mathsf{val}(\rho \cup g, \mathcal{W}_\ell)]$ and the variance of the satisfied weight $\mathbf{Var}_g[\mathsf{val}(\rho \cup g, \mathcal{W}_\ell)]$. We can thus classify instances as concentrated or non-concentrated. If more than t variables in S are labelled by instance ℓ (where $t = O_{k,\varepsilon}(1)$ is some parameter to be chosen), we call instance ℓ saturated. If every unsaturated instance is concentrated, then we are done with this S and ρ, and this branch of the tree gets terminated.

Otherwise, we know that there some unsaturated instance ℓ which is not concentrated. We know that this instance ℓ must have some variable $x \in V \backslash S$ which has high *active degree* (this is the degree after taking into account the partial assignment ρ). The algorithm now takes two cases:

- **Case 1:** If this high-active-degree variable x is involved in a high-weight constraint on $\{x, y\}$ for some $y \in V \backslash S$, then we include both x, y into the set S, and consider all 4 possible extensions of ρ to this new S. x, y are both labelled with instance ℓ.

- **Case 2:** Otherwise, every constraint involving x is low-weight (and in particular there must be many of them), and in this case we include x into the set S, and consider both possible extensions of ρ to this new S. x is labelled with instance ℓ.

This concludes the first stage of the algorithm, which created a tree whose leaves contain various (S, ρ) pairs.

For the second stage of the algorithm we visit each leaf (S, ρ). We choose a uniformly random $g : V \backslash S \to \{0,1\}$, and consider for every $h : S \to \{0,1\}$, the assignment $h \cup g : V \to \{0,1\}$. Note that we go over all assignments to the set S, independent of the partial assignment to S associated with the leaf.

The analysis: At the end of the evolution, at every leaf of the tree every instance is either highly-concentrated or saturated. If instance ℓ is highly-concentrated, we will have the property that the random assignment to $V \backslash S$ has the right approximation factor for instance ℓ. If the instance ℓ is saturated, then we know that there are many variables in S labelled by instance ℓ; and at the time these variables were brought into S, they had high active degree.

The main part of the analysis is then a delicate *perturbation* procedure, which starts with the partial assignment $h^* \stackrel{\text{def}}{=} f^*|_S$, and perturbs it to some $h' : S \to \{0,1\}$ with a certain robustness property. Specifically, it ensures that for every saturated instance $\ell \in [k]$. we have $\mathsf{val}(h' \cup g, \mathcal{W}_\ell)$ is

at least as large as the total weight in instance ℓ of all constraints not wholly contained within S. At the same time, the perturbation ensures that for unsaturated instances $\ell \in [k]$, $\mathsf{val}(h' \cup g, \mathcal{W}_\ell)$ is almost as large as $\mathsf{val}(h^* \cup g, \mathcal{W}_\ell)$. This yields the desired Pareto approximation. The perturbation procedure modifies the assignment h^* at a few carefully chosen variables (at most two variables per saturated instance). After picking the variables for an instance, if the variables were brought into S by Case 1, we can satisfy the heavy constraint involving them. Otherwise, we use a Lipschitz concentration bound to argue that a large fraction of the constraints involving the variable and $V \backslash S$ can be satisfied; this is the second place where we use the randomness in the choice of g.

As we mentioned earlier, this perturbation is necessary! It is not true the assignment $h^* \cup g$ will give a good Pareto approximation with good probability [4].

Improved approximation, and generalization: To get the claimed $(\frac{1}{2} - \varepsilon)$-Pareto approximation for the $q = w = 2$ case, we replace the uniformly random choice of $g : V \backslash S \to \{0,1\}$ by a suitable LP relaxation + randomized rounding strategy. Concretely, at every leaf (S, ρ), we do the following. First we write an LP relaxation of the residual MAX-2-CSP problem. Then, using a rounding algorithm of Trevisan (which has some desirable smoothness properties), we choose $g : V \backslash S \to \{0,1\}$ by independently rounding each variable. Finally, for all $h : S \to \{0,1\}$, we consider the assignment $h \cup g$. The analysis is nearly identical (but crucially uses the smoothness of the rounding), and the improved approximation comes from the improved approximation factor of the classical LP relaxation for MAX-2-CSP.

The generalization of this algorithm to general q, w is technical but straightforward. One notable change is that instead of taking 2 cases each time we grow the tree, we end up taking w cases. In case j, we have a set of j variables such that the total weight of constraints involving all the j variables is large, however for every remaining variable z, the weight of contraints involving all the j variables together with z is small. The analysis of the perturbation is similar.

The algorithm for Max-w-SAT uses the fact that the LP rounding gives a $3/4$ approximation for Max-w-SAT. Moreover, since a Max-w-SAT constraint can be satisfied by perturbing any one variable, the algorithm does not require a tree of evolutions. It only maintains a set of "influential" variables, and hence, is simpler.

1.5 Related Work

The theory of *exact* multiobjective optimization has been very well studied, (see eg. [PY00, Dia11] and the references therein).

The only directly comparable work for simultaneous approximation algorithms for CSPs we are aware of is the work of Glaßer *et al.* [GRW11] [5]. They give a $1/2$-Pareto approximation for Max-SAT with a running time of $n^{O(k^2)}$. For bounded width clauses, our algorithm does better in both approximation guarantee and running time.

For Max-CUT, there are a few results of a similar flavor. For two graphs, the results of Angel *et al.* [ABG06] imply a 0.439-Pareto approximation algorithm (though their actual results are incomparable to ours). Bollobás and Scott [BS04] asked what is the largest simultaneous cut in two unweighted graphs with m edges each. Kuhn and Osthus [KO07], using the second moment method, proved that for k simultaneous unweighted instances, there is a simultaneous cut that cuts

[4] See Section E for an example

[5] They also give Pareto approximation results for simultaneous TSP (also see references therein).

at least $m/2 - O(\sqrt{km})$ edges in each instance, and give a deterministic algorithm to find it (this leads to a $(\frac{1}{2} - o(1))$-Pareto approximation for unweighted instances with sufficiently many edges). Our main theorem implies the same Pareto approximation factor for simultaneous MAX-CUT on general weighted instances, while for k-fold simultaneous MAX-CUT on unweighted instances, our Theorem 1.5 gives a $\left(\frac{1}{2} + \Omega(\frac{1}{k^2})\right)$-minimum approximation algorithm.

1.6 Discussion

We have only made initial progress on what we believe is a large number of interesting problems in the realm of simultaneous approximation of CSPs. We list here a few of the interesting directions for further research:

1. When designing SDP-based algorithms for the classical MAX-CSP problems, we are usually only interested in the expected value of the rounded solution. For k-fold simultaneous MAX-$\mathcal{F}$-CSP with $k > 1$, we are naturally led to the question of how concentrated the value of the solution output by the rounding is around its mean.

   Decorrelation of SDP rounding arises in recent algorithms [BRS11, RT12, GS11] based on SDP hierarchies. It would be interesting to see if such ideas could be useful in this context.

   Another interesting question of this flavor is whether there are natural conditions under which the Goemans-Williamson hyperplane rounding gives a good solution for MAXCUT with high probability.

2. When $k = O(1)$, for each $\mathcal{F}$, one can ask the question: what is the best Pareto approximation factor achievable for k-fold MAX-$\mathcal{F}$-CSP in polynomial time? While in Theorem 1.2 we do not focus on giving improved approximation factors for special $\mathcal{F}$, our methods will give better approximation factors for any $\mathcal{F}$ which has a good LP relaxation that comes equipped with a sufficiently smooth independent-rounding algorithm.

   It would be very interesting if one could employ SDPs for approximating simultaneous MAX-$\mathcal{F}$-CSP . A particularly nice question here: *Is there a polynomial time* 0.878*-Pareto approximation algorithm for* $O(1)$*-fold simultaneous* MAX-CUT*?* We do not even know a $(1/2 + \varepsilon)$-Pareto approximation algorithm (but note that Theorem 1.5 does give this for $O(1)$-fold simultaneous *unweighted* MAX-CUT).

3. As demonstrated by hardness result for MAX-w-SAT given in Proposition 1.1, even for constant k, the achievable approximation factor can be strictly smaller than its classical counterpart. It would be very interesting to have a systematic theory of hardness reductions for simultaneous CSPs for $k = O(1)$. The usual paradigm for proving hardness of approximation based on label cover and long codes seems to break down completely for simultaneous CSPs.

1.7 Organization of this paper

We first present the notation required for our algorithms in Section 2. We then describe our Pareto approximation algorithm for MAX-2-AND (which is equivalent to MAX-2-CSP$_2$), and its generalization to MAX-w-CSP$_q$ in Sections 3 and 4 respectively. We then present our improved Pareto approximation for MAX-w-SAT in Section 5.

We present the additional results in the appendix. The dichotomy theorem for the hardness of arbitrary CSPs is presented in Section A, followed by our improved minimum approximation algorithm for unweighted MAX-CUT in Section B, and the SDP integrality gaps in Section C.

2 Notation for the main algorithms

We now define some common notation that will be required for the following sections on algorithms for MAX-2-AND and and for general MAX-$\mathcal{F}$-CSP. For the latter, will stop referring to the set of predicates $\mathcal{F}$, and simply present an algorithm for the problem MAX-w-CSP$_q$: this is the MAX-$\mathcal{F}$-CSP problem, where $\mathcal{F}$ equals the set of all predicates on w variables from the domain $[q]$. For MAX-2-AND, the alphabet q and arity w are both 2.

Let V be a set of n variables. Each variable will take values from the domain $[q]$. Let $\mathcal{C}$ denote a set of constraints of interest on V (for example, for studying MAX-2-AND, $\mathcal{C}$ would be the set of AND constraints on pairs of literals of variables coming from V). We use the notation $v \in C$ to denote that the v is one of the variables that the constraint C depends on. Analogously, we denote $T \subseteq C$ if C depends on all the variables in T. A weighted MAXCSP instance on V is given by a weight function $\mathcal{W} : \mathcal{C} \to \mathbb{R}_+$, where for $C \in \mathcal{C}$, $\mathcal{W}(C)$ is the weight of the constraint C. We will assume that $\sum_{C \in \mathcal{C}} \mathcal{W}(C) = 1$.

A partial assignment ρ is a pair (S_ρ, h_ρ), where $S_\rho \subseteq V$ and $h_\rho : S_\rho \to [q]$. (We also call a function $h : S \to [q]$, a partial assignment, when S is understood from the context). We say a contraint $C \in \mathcal{C}$ is active given ρ if C depends on some variable in $V \backslash S_\rho$, and there exists full assignments $g_0, g_1 : V \to [q]$ with $g_i|_{S_\rho} = h_\rho$, such that C evaluates to FALSE under the assignment g_0 and C evaluates to TRUE under the assignment g_1. (colloquially: C's value is not fixed by ρ). We denote by $\mathsf{Active}(\rho)$ the set of constraints from $\mathcal{C}$ which are active given ρ. For a partial assignment ρ and $C \in \mathcal{C} \backslash \mathsf{Active}(\rho)$, let $C(\rho) = 1$ if C's value is fixed to TRUE by ρ, and let $C(\rho) = 0$ if C's value is fixed to FALSE by ρ. For disjoint subsets $S_1, S_2 \subseteq V$ and partial assignments $f_1 : S_1 \to [q]$ and $f_2 : S_2 \to [q]$, let $f = f_1 \cup f_2$ denote the assignment $f : S_1 \cup S_2 \to [q]$ with $f(x) = f_1(x)$ if $x \in S_1$, and $f(x) = f_2(x)$ if $x \in S_2$. Abusing notation, for a partial assignment ρ and an assignment $g : V \backslash S_{\rho^\star} \to [q]$, we often write $\rho \cup g$ instead of $h_\rho \cup g$. For two constraints $C_1, C_2 \in \mathcal{C}$, we say $C_1 \sim_\rho C_2$ if they share a variable that is contained in $V \backslash S_\rho$.

Define the *active degree given* ρ of a variable $v \in V \backslash S_\rho$ by:

$$\mathsf{activedegree}_\rho(v, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{C \in \mathsf{Active}(\rho), C \ni v} \mathcal{W}(C).$$

For a subset $T \subseteq V \backslash S_\rho$ of variables, define its active degree given ρ by:

$$\mathsf{activedegree}_\rho(T, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{C \in \mathsf{Active}(\rho), C \supseteq T} \mathcal{W}(C).$$

Define the active degree of the whole instance $\mathcal{W}$ given ρ:

$$\mathsf{activedegree}_\rho(\mathcal{W}) \stackrel{\text{def}}{=} \sum_{v \in V \backslash S_\rho} \mathsf{activedegree}_\rho(v, \mathcal{W}).$$

For a partial assignment ρ, we define its value on an instance $\mathcal{W}$ by:

$$\mathsf{val}(\rho, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{C \in \mathcal{C} \backslash \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho).$$

Thus, for a total assignment $f : V \to [q]$ extending ρ, we have the equality:

$$\mathsf{val}(f, \mathcal{W}) - \mathsf{val}(\rho, \mathcal{W}) = \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(f).$$

3 Simultaneous Max-2-AND

In this section, we give our approximation algorithm for simultaneous Max-2-AND. Via a simple reduction given Section 4.1, this implies the $q = w = 2$ case of our main theorem, Theorem 1.2.

3.1 Random Assignments

We begin by giving a sufficient condition for the value of a Max-2-AND to be highly concentrated under independent random assignments to the variables.

Let ρ be a partial assignment. Let $p : V \backslash S_\rho \to [0, 1]$ be such that $p(v) \in [\frac{1}{4}, \frac{3}{4}]$ for each $v \in V \backslash S_\rho$. Let $g : V \backslash S_\rho \to [q]$ be a random assignment obtained by sampling $g(v)$ for each v independently with $\mathbf{E}[g(v)] = p(v)$. Define the random variable

$$Y \stackrel{\text{def}}{=} \mathsf{val}(\rho \cup g, \mathcal{W}) - \mathsf{val}(\rho, \mathcal{W}) = \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(g).$$

The random variable Y measures the contribution of active constraints to $\mathsf{val}(\rho \cup g, \mathcal{W})$. Note that the two quantities $\mathbf{E}[Y]$ and $\mathbf{Var}[Y]$ can be computed efficiently given p. We denote these by $\mathsf{TrueMean}_\rho(p, \mathcal{W})$ and $\mathsf{TrueVar}_\rho(p, \mathcal{W})$. The following lemma proves that either Y is concentrated, or there exists an active variable that contributes a significant fraction of the total active-degree of the instance.

Lemma 3.1 *Let p, Y be as above.*

1. *If $\mathsf{TrueVar}_\rho(p, \mathcal{W}) < \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_\rho(p, \mathcal{W})^2$ then $\Pr[Y < (1 - \varepsilon_0)\, \mathbf{E}[Y]] < \delta_0$.*
2. *If $\mathsf{TrueVar}_\rho(p, \mathcal{W}) \geqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_\rho(p, \mathcal{W})^2$, then there exists $v \in V \backslash S_\rho$ such that*

$$\mathsf{activedegree}_\rho(v, \mathcal{W}) \geqslant \frac{\varepsilon_0^2 \delta_0}{64} \cdot \mathsf{activedegree}_\rho(\mathcal{W}).$$

The above lemma is a special case of Lemma 4.2 which is proved in Section 4.2, and hence we skip the proof. The first part is then a simple application of the Chebyshev inequality. For the second part, we use the assumption that $\mathsf{TrueVar}$ is large, to deduce that there exists a constraint C such that the total weight of constraints that share a variable from $V \backslash S$ with C, *i.e.*, $\sum_{C_2 \sim_S C} \mathcal{W}(C_2)$, is large. It then follows that at least one variable $v \in C$ must have large activedegree given S.

3.2 LP relaxations

Let $(c_\ell)_{\ell \in [k]}$ be the given target values for the Pareto approximation problem. Given a partial assignment ρ, we can write the feasibility linear program for simultaneous Max-2-AND as shown

in figure 1, 2. In this LP, for a constraint C, C^+ (C^-) denotes set of variables that appears as a positive (negative) literal in C.

For $\vec{t}$, $\vec{z}$ satisfying linear constraints $\mathsf{MAX2AND\text{-}LP}_1(\rho)$, let $\mathsf{smooth}(\vec{t})$ denote the map $p : V\backslash S_\rho \to [0,1]$ with $p(v) = \frac{1}{4} + \frac{t_v}{2}$. Note that $p(v) \in [1/4, 3/4]$ for all v.

Given $\vec{t}, \vec{z}$ satisfying $\mathsf{MAX2AND\text{-}LP}_1$, the rounding algorithm from [Tre98] samples each variable v independently with probabily $\mathsf{smooth}(\vec{t})(v)$. Note that this rounding algorithm is *smooth* in the sense that each variable is sampled independently with a probability that is bounded away from 0 and 1. This will be crucial for our algorithm. The following theorem from [Tre98] proves that this rounding algorithm finds a good integral assignment.

Lemma 3.2 ([Tre98]) *Let ρ be a partial assignment.*

1. **Relaxation:** *For every $g_0 : V\backslash S_\rho \to \{0,1\}$, there exist $\vec{t}$, $\vec{z}$ satisfying $\mathsf{MAX2AND\text{-}LP}_1(\rho)$ such that for every* Max-2-AND *instance $\mathcal{W}$:*
$$\sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C = \mathsf{val}(g_0 \cup \rho, \mathcal{W}).$$

2. **Rounding:** *Suppose $\vec{t}, \vec{z}$ satisfy $\mathsf{MAX2AND\text{-}LP}_1(\rho)$. Then for every* Max-2-AND *instance $\mathcal{W}$:*
$$\mathsf{val}(\rho, \mathcal{W}) + \mathsf{TrueMean}_\rho(\mathsf{smooth}(\vec{t}), \mathcal{W}) \geqslant \frac{1}{2} \cdot \sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C.$$

Proof: We begin with the first part. For $v \in S_\rho$, define $t_v = \rho(v)$.. For $v \in V\backslash S_\rho$, define $t_v = g_0(v)$. For $C \in \mathcal{C}$, define $z_C = 1$ if $C(g_0 \cup \rho) = 1$, and define $z_C = 0$ otherwise. It is easy to see that these $\vec{t}, \vec{z}$ satisfies $\mathsf{MAX2AND\text{-}LP}_1(\rho)$, and that for every instance $\mathcal{W}$:

$$\sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C = \mathsf{val}(g_0 \cup \rho, W).$$

Now we consider the second part. Let $\mathcal{W}$ be any instance of Max-2-AND. Let $p = \mathsf{smooth}(\vec{t})$. Let $g : V\backslash S_\rho \to \{0,1\}$ be sampled as follows: independently for each $v \in V\backslash S_\rho$, $g(v)$ is sampled from $\{0,1\}$ such that $\mathbf{E}[g(v)] = p(v)$. We have:

$$\mathsf{val}(\rho, \mathcal{W}) + \mathsf{TrueMean}_\rho(\mathsf{smooth}(\vec{t}), \mathcal{W}) = \sum_{C \in \mathcal{C}\backslash \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho) + \mathbf{E}\left[\sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho \cup g) \right]. \tag{1}$$

We will now understand the two terms of the right hand side.

For $C \in \mathcal{C}\backslash\mathsf{Active}(\rho)$, it is easy to verify that if $z_C > 0$, we must have $C(\rho) = 1$. Thus:

$$\sum_{C \in \mathcal{C}\backslash \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho) \geqslant \sum_{C \in \mathcal{C}\backslash \mathsf{Active}(\rho)} \mathcal{W}(C) z_C.$$

To understand the second term, we have the following claim.

Claim 3.3 *For $C \in \mathsf{Active}(\rho)$, $\mathbf{E}[C(\rho \cup g)] \geqslant \frac{1}{2} \cdot z_C$.*

$$\begin{array}{rcll} z_C & \leqslant & t_v & \forall C \in \mathcal{C}, v \in C^+ \\ z_C & \leqslant & 1 - t_v & \forall C \in \mathcal{C}, v \in C^- \\ 1 \geqslant t_v & \geqslant & 0 & \forall v \in V \backslash S_\rho \\ t_v & = & h_\rho(v) & \forall v \in S_\rho \end{array}$$

Figure 1: Linear inequalities $\mathsf{MAX2AND\text{-}LP}_1(\rho)$

$$\begin{array}{c} \sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) \cdot z_C \geqslant c_\ell \quad \forall \ell \in [k] \\ \vec{t}, \vec{z} \text{ satisfy } \mathsf{MAX2AND\text{-}LP}_1(\rho). \end{array}$$

Figure 2: Linear inequalities $\mathsf{MAX2AND\text{-}LP}_2(\rho)$

Proof: Suppose there are exactly h variables in C which are not in S_ρ. We have $h \leqslant 2$.

$$\begin{aligned} \mathbf{E}[C(\rho \cup g)] = \Pr[C \text{ is satisfied by } \rho \cup g] &= \left(\prod_{v \in C^+, v \in V \backslash S_\rho} \frac{1}{4} + \frac{t_v}{2} \right) \cdot \left(\prod_{v \in C^-, v \in V \backslash S_\rho} \frac{1}{4} + \frac{1 - t_v}{2} \right) \\ &\geqslant \left(\frac{1}{4} + \frac{z_C}{2} \right)^h \geqslant \left(\frac{1}{4} + \frac{z_C}{2} \right)^2 \geqslant \frac{z_C}{2}. \end{aligned}$$

■

This claim implies that:

$$\mathbf{E}\left[\sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho \cup g) \right] \geqslant \frac{1}{2} \sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C.$$

Substituting back into Equation (1), we get the Lemma. ■

3.3 The Algorithm

We now give our Pareto approximation algorithm for Max-2-AND in Figure 3.

Input: k instances of Max-2-AND $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on the variable set V, $\varepsilon > 0$ and target objective values $c_1, \ldots, c_k$.
Output: An assignment to V
Parameters: $\delta_0 = \frac{1}{10(k+1)}$, $\varepsilon_0 = \varepsilon$, $\gamma = \frac{\varepsilon_0^2 \delta_0}{16}$, $t = \left\lceil \frac{20k^2}{\gamma} \log \frac{k}{\gamma} \right\rceil$

1. Initialize tree T to be an empty quaternary tree (i.e., just 1 root node). Nodes of the tree will be indexed by strings in $(\{0,1\}^2)^*$.
2. With each node ν of the tree, we associate:
   (a) A partial assignment ρ_ν.
   (b) A special pair of variables $\mathcal{T}_\nu^1, \mathcal{T}_\nu^2 \in V \backslash S_{\rho_\nu}$.
   (c) A special instance $\mathcal{I}_\nu \in [k]$.
   (d) A collection of integers $\mathsf{count}_{\nu,1}, \ldots, \mathsf{count}_{\nu,k}$.
   (e) A trit representing whether the node ν is living, dead, or exhausted.

3. Initialize the root node ν_0 to (1) $\rho_{\nu_0} \leftarrow (\varnothing, \varnothing)$, (2) $\forall \ell \in [k], \mathsf{count}_{\nu_0,\ell} \leftarrow 0$, (3) living.
4. While there is a living leaf ν of T, do the following:
   (a) Check the feasibility of linear inequalities $\mathsf{MAX2AND\text{-}LP}_2(\rho_\nu)$.
      i. If there is a feasible solution $\vec{t}, \vec{z}$, then define $p_\nu : V \backslash S_{\rho_\nu} \to [0,1]$ as $p_\nu = \mathsf{smooth}(\vec{t})$.
      ii. If not, then declare ν to be dead and return to Step 4.
   (b) For each $\ell \in [k]$, compute $\mathsf{TrueVar}_{\rho_\nu}(p_\nu, \mathcal{W}_\ell)$ and $\mathsf{TrueMean}_{\rho_\nu}(p_\nu, \mathcal{W}_\ell)$.
   (c) If $\mathsf{TrueVar}_{\rho_\nu}(p_\nu, \mathcal{W}_\ell) \geqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_{\rho_\nu}(p_\nu, \mathcal{W}_\ell)^2$, then set $\mathsf{flag}_\ell \leftarrow$ TRUE, else set $\mathsf{flag}_\ell \leftarrow$ FALSE.
   (d) Choose the smallest $\ell \in [k]$, such that $\mathsf{count}_\ell < t$ AND $\mathsf{flag}_\ell =$ TRUE (if any):
      i. Find $x \in V \backslash S_{\rho_\nu}$ that maximizes $\mathsf{activedegree}_{\rho_\nu}(x, \mathcal{W}_\ell)$. Note that it will satisfy $\mathsf{activedegree}_{\rho_\nu}(x, \mathcal{W}_\ell) \geqslant \gamma \cdot \mathsf{activedegree}_{\rho_\nu}(\mathcal{W}_\ell)$.
      ii. Among all the active constraints $C \in \mathcal{C}$ such that $x \in C$ and $C \cap S_{\rho_\nu} = \varnothing$, find the one that maximizes $\mathcal{W}_\ell(C)$. Call this constraint $C^\star$. Let y be the other variable contained in $C^\star$ (if there is no other variable, set $y = x$).
      Set $\mathcal{T}_\nu^1 \leftarrow x$ and $\mathcal{T}_\nu^2 \leftarrow y$. Set $\mathcal{I}_\nu \leftarrow \ell$.
      iii. Create four children of ν, with labels $\nu b_1 b_2$ for each $b_1, b_2 \in \{0,1\}$ and set
         - $\rho_{\nu b_1 b_2} \leftarrow (S_{\rho_\nu} \cup \{\mathcal{T}_\nu^1, \mathcal{T}_\nu^2\}, h^{b_1 b_2})$, where $h^{b_1 b_2}$ extends h_{ρ_ν} by $h^{b_1 b_2}(\mathcal{T}_\nu^1) = b_1$ and $h^{b_1 b_2}(\mathcal{T}_\nu^2) = b_2$.
         - $\forall \ell' \in [k]$ with $\ell' \neq \ell$, set $\mathsf{count}_{\nu b_1 b_2, \ell'} \leftarrow \mathsf{count}_{\nu, \ell'}$. Set $\mathsf{count}_{\nu b_1 b_2, \ell} \leftarrow \mathsf{count}_{\nu,\ell} + 1$.
         - Set $\nu b_1 b_2$ to be living.
   (e) If no such ℓ exists, declare ν to be exhausted.
5. Now every leaf of T is either exhausted or dead. For each exhausted leaf ν of T:
   (a) Let $g_\nu : V \backslash S_{\rho_\nu} \to \{0,1\}$ be a random assignment where, for each $v \in V \backslash S_{\rho_\nu}$, $g_\nu(v)$ is sampled independently with $\mathbf{E}[g_\nu(v)] = p_\nu(v)$. Note that $\mathbf{E}[g_\nu(v)] \in [\frac{1}{4}, \frac{3}{4}]$.
   (b) For every assignment $h : S_{\rho_\nu} \to \{0,1\}$, compute $\mathsf{out}_{h,g_\nu} \leftarrow \min_{\ell \in [k]} \frac{\mathsf{val}(h \cup g_\nu, \mathcal{W}_\ell)}{c_\ell}$. If $c_\ell = 0$ for some $\ell \in [k]$, we interpret $\frac{\mathsf{val}(h \cup g_\nu, \mathcal{W}_l)}{c_\ell}$ as $+\infty$.
6. Output the largest out_{h,g_ν} seen, and the assignment $h \cup g_\nu$ that produced it.

Figure 3: Algorithm SIM-MAX2AND for approximating weighted simultaneous MAX-2-AND

3.4 Analysis

Notice that the depth of the tree T is at most kt, and that for every ν, we have that $|S_{\rho_\nu}| \leqslant 2kt$. This implies that the running time is at most $2^{O(kt)} \cdot \mathrm{poly}(n)$.

Let $f^\star : V \to \{0,1\}$ be an assignment such that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for each $\ell \in [k]$. Let $\nu^\star$ be the the unique leaf of the tree T for which $\rho_{\nu^\star}$ is consistent with $f^\star$. (This $\nu^\star$ can be found as follows: start with ν equal to the root. Set ν to equal the unique child of ν for which ρ_ν is consistent with $f^\star$, and repeat until ν becomes a leaf. This leaf is $\nu^\star$). Observe that since $f^\star$ is an assignment such that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for every $\ell \in [k]$, by picking $g_0 = f^\star|_{V \setminus S^\star}$ in part 1 of Lemma 3.2, we know that $\mathsf{MAX2AND\text{-}LP}_2(\rho^\star)$ is feasible, and hence $\nu^\star$ must be an exhausted leaf (and not dead).

Define $\rho^\star = \rho_{\nu^\star}$, $S^\star = S_{\rho^\star}$, $h^\star = h_{\rho^\star}$, and $p^\star = p_{\nu^\star}$. At the completion of Step 4, if $\ell \in [k]$ satisfies $\mathsf{count}_{\nu^\star,\ell} = t$, we call instance ℓ a *high variance* instance. Otherwise we call instance ℓ a *low variance* instance.

3.4.1 Low Variance Instances

First we show that for the leaf ν^* in Step 5, combining the partial assignment $h^\star$ with a random assignment $g_{\nu^\star}$ in step 5(b) is good for any low variance instances with high probability.

Lemma 3.4 *Let $\ell \in [k]$ be any low variance instance. For the leaf node $\nu^\star$, let $g_{\nu^\star}$ be the random assignment sampled in Step 5.(a). of* Sim-Max2AND. *Then with probability at least $1 - \delta_0$, the assignment $f = h^\star \cup g_{\nu^\star}$ satisfies:*

$$\Pr_{g_{\nu^\star}}\left[\mathsf{val}(f, \mathcal{W}_\ell) \geqslant (1/2 - \varepsilon/2) \cdot c_\ell\right] \geqslant 1 - \delta_0.$$

Proof: For every low variance instance ℓ, we have that $\mathsf{TrueVar}_{\rho_{\nu^\star}}(p^\star, \mathcal{W}_\ell) < \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_{\rho_{\nu^\star}}(p^\star, \mathcal{W}_\ell)^2$. Define $Y \stackrel{\text{def}}{=} \mathsf{val}(\rho^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) - \mathsf{val}(\rho^\star, \mathcal{W}_\ell)$. By Lemma 3.1, we have $\Pr[Y < (1-\varepsilon_0)\mathbf{E}[Y]] < \delta_0$. Thus, with probability at least $1 - \delta_0$, we have,

$$\begin{aligned}
\mathsf{val}(f, \mathcal{W}_\ell) &\geqslant \mathsf{val}(\rho^\star, \mathcal{W}_\ell) + (1-\varepsilon_0)\mathbf{E}[Y] \\
&= \mathsf{val}(\rho^\star, \mathcal{W}_\ell) + (1-\varepsilon_0) \cdot \mathsf{TrueMean}_{\rho^\star}(\mathsf{smooth}(\vec{t}), \mathcal{W}_\ell) \\
&= (1-\varepsilon_0) \cdot \left(\mathsf{val}(\rho^\star, \mathcal{W}_\ell) + \mathsf{TrueMean}_{\rho^\star}(\mathsf{smooth}(\vec{t}), \mathcal{W}_\ell)\right) \\
&\geqslant \frac{1}{2} \cdot (1-\varepsilon_0) \cdot \sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) \cdot z_C \geqslant \frac{1}{2} \cdot (1-\varepsilon_0) \cdot c_\ell \geqslant \left(\frac{1}{2} - \frac{\varepsilon}{2}\right) \cdot c_\ell,
\end{aligned}$$

where we have used the second part of Lemma 3.2. ■

Next, we will consider a small perturbation of $h^\star$ which will ensure that the algorithm performs well on high variance instances too. We will ensure that this perturbation does not affect the success on the low variance instances.

3.4.2 High variance instances

Fix a high variance instance ℓ. Let ν be an ancestor of $\nu^\star$ with $\mathcal{I}_\nu = \ell$. Define:

$$\mathsf{activedegree}_\nu \stackrel{\text{def}}{=} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu^1, \mathcal{W}_\ell).$$

Let $\mathcal{C}_\nu$ be the set of all constraints C containing $\mathcal{T}_\nu^1$ which are active given ρ_ν. We call a constraint C in $\mathcal{C}_\nu$ a backward constraint if C only involves variables from $S_{\rho_\nu} \cup \{\mathcal{T}_\nu^1\}$. Otherwise we call C in $\mathcal{C}_\nu$ a forward constraint. Let $\mathcal{C}_\nu^{\mathsf{backward}}$ and $\mathcal{C}_\nu^{\mathsf{forward}}$ denote the sets of these constraints. Finally, we denote $\mathcal{C}_\nu^{\mathsf{out}}$ the set of binary constraints on $\mathcal{T}_\nu^1$ and a variable from $V \backslash S^\star$.

Define backward degree and forward degree as follows:

$$\mathsf{backward}_\nu \stackrel{\text{def}}{=} \sum_{C \in \mathcal{C}_\nu^{\mathsf{backward}}} \mathcal{W}_\ell(C),$$

$$\mathsf{forward}_\nu \stackrel{\text{def}}{=} \sum_{C \in \mathcal{C}_\nu^{\mathsf{forward}}} \mathcal{W}_\ell(C).$$

Note that:

$$\mathsf{activedegree}_\nu = \mathsf{backward}_\nu + \mathsf{forward}_\nu.$$

Now we consider variable $\mathcal{T}_\nu^2$. Let $\mathsf{heaviest}_\nu$ be the total $\mathcal{W}_\ell$ weight of all the constraints containing both $\mathcal{T}_\nu^1$ and $\mathcal{T}_\nu^2$. Based on all this, we classify ν into one of three categories:

1. If $\mathsf{backward}_\nu \geqslant \frac{1}{2} \cdot \mathsf{activedegree}_\nu$, then we call ν a typeA node.
2. Otherwise, if $\mathsf{heaviest}_\nu \geqslant \frac{1}{100tk} \cdot \mathsf{activedegree}_\nu$, then we call ν a typeB node. In this case we have some $\mathcal{W}_\ell$ constraint C containing $\mathcal{T}_\nu^1$ and $\mathcal{T}_\nu^2$ with $\mathcal{W}_\ell(C) \geqslant \frac{1}{1600tk} \cdot \mathsf{activedegree}_\nu$.
3. Otherwise, we call ν a typeC node. In this case, for every $v \in V \backslash S_{\rho_\nu}$, the total weight of the constraints involving v and $\mathcal{T}_\nu^1$, *i.e.*, $\mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu^1 \cup v, \mathcal{W}_\ell)$ is bounded by $\frac{1}{100tk} \cdot \mathsf{activedegree}_\nu$. In particular, every constraint $C \in \mathcal{C}_\nu^{\mathsf{forward}}$ must have $\mathcal{W}_\ell(C) < \frac{1}{100tk} \cdot \mathsf{activedegree}_\nu$. Since $|S^\star| \leqslant 2tk$, the total weight of constraints containing $\mathcal{T}_\nu^1$ and some variable in $S^\star \backslash S_{\rho_\nu}$ is at most $|S^\star \backslash S_{\rho_\nu}| \cdot \frac{1}{100tk} \cdot \mathsf{activedegree}_\nu$ which is at most $\frac{2}{100} \cdot \mathsf{activedegree}_\nu$. Hence we have:

$$\begin{aligned}\sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) &= \mathsf{forward}_\nu - \left\{ \begin{smallmatrix} \text{total weight of constraints containing} \\ \mathcal{T}_\nu^1 \text{ and some variable in } S^\star \backslash S_{\rho_\nu} \end{smallmatrix} \right\} \\ &\geqslant \left(\frac{1}{2} - \frac{2}{100} \right) \mathsf{activedegree}_\nu > \frac{1}{4} \cdot \mathsf{activedegree}_\nu.\end{aligned}$$

For nodes ν which are typeC, the variable $\mathcal{T}_\nu^1$ has a large fraction of its active degree coming from constraints between $\mathcal{T}_\nu^1$ and $V \backslash S^\star$.

For a partial assignment $g : V \backslash S^\star \to \{0,1\}$, we say that g is Cgood for ν if there exists a setting of variable $\mathcal{T}_\nu^1$ that satisfies at least $\frac{1}{64} \cdot \mathsf{activedegree}_\nu$ weight amongst constraints containing variable $\mathcal{T}_\nu^1$ and some other variable in $V \backslash S^\star$. The next lemma shows that for every typeC node ν, with high probability, the random assignment $g_{\nu^\star} : V \backslash S^\star \to \{0,1\}$ is Cgood for ν.

Lemma 3.5 *Consider a* typeC *node ν. Suppose $g : V \backslash S^\star \to \{0,1\}$ is a partial assignment obtained by independently sampling $g(v)$ with $\mathbf{E}[g(v)] \in [1/4, 3/4]$ for each $v \in V \backslash S^\star$. Then:*

$$\Pr_g[g \text{ is Cgood for } \nu] \geqslant 1 - 2 \cdot e^{-tk/100}.$$

Proof: Let $\ell = \mathcal{I}_\nu$.

For each constraint $C \in \mathcal{C}_\nu^{\mathsf{out}}$ and each $g : \{0,1\}^{V\setminus S^\star} \to \{0,1\}$, define $Z_C^{(1)}(g), Z_C^{(0)}(g) \in \{0,1\}$ as follows. $Z_C^{(1)}(g)$ equals 1 iff C is satisfied by extending the assignment g with $\mathcal{T}_\nu^1 \leftarrow 1$. Similarly, $Z_C^{(0)}(g)$ equals 1 iff C is satisfied by extending the assignment g with $\mathcal{T}_\nu^1 \leftarrow 0$.

For $b = 0, 1$, we define $\mathsf{score}^{(b)} : \{0,1\}^{V\setminus S^\star} \to \mathbb{R}$ as follows:

$$\mathsf{score}^{(b)}(g) \stackrel{\text{def}}{=} \sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) \cdot Z_C^{(b)}(g).$$

In words, $\mathsf{score}^{(b)}(g)$ is the total weight of constraints between $\mathcal{T}_\nu^1$ and $V\setminus S^*$ satisfied by setting $\mathcal{T}_\nu^1$ to b and setting $V\setminus S^*$ according to g.

Note that since $g(v)$ is sampled independently for $v \in V\setminus S^\star$ with $\mathbf{E}[g(v)] \in [1/4, 3/4]$, we have $\mathbf{E}_g[Z_C^{(1)}(g) + Z_C^{(0)}(g)] \geqslant \frac{1}{4}$. Thus:

$$\begin{aligned}\mathbf{E}_g[\mathsf{score}^{(1)}(g) + \mathsf{score}^{(0)}(g)] &= \sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} W_\ell(C)\mathbf{E}[Z_C^{(1)}(g)] + \sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} W_\ell(C)\mathbf{E}[Z_C^{(0)}(g)] \\ &\geqslant \frac{1}{4} \sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C).\end{aligned}$$

So one of $\mathbf{E}[\mathsf{score}^{(1)}(g)]$ and $\mathbf{E}[\mathsf{score}^{(0)}(g)]$ is at least $\frac{1}{8}\sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) \geqslant \frac{1}{32}\mathsf{activedegree}_\nu$. Suppose it is $\mathbf{E}[\mathsf{score}^{(1)}(g)]$ (the other case is identical). We are going to use McDiarmid's inequality to show the concentration of $\mathsf{score}^{(1)}(g)$ around its mean[6].

Since ν is typeC, we know that for every vertex $v \in V\setminus S^\star$, changing g on just v can change the value of $\mathsf{score}^{(1)}(g)$ by at most $c_v \stackrel{\text{def}}{=} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu^1 \cup v, \mathcal{W}_\ell) \leqslant \frac{1}{100tk} \cdot \mathsf{activedegree}_\nu$. Thus by McDiarmid's inequality (Lemma D.1),

$$\begin{aligned}\Pr_g[g \text{ is not } \mathsf{Cgood} \text{ for } \nu] &\leqslant \Pr_g\left[\mathsf{score}^{(1)}(g) < \frac{1}{64} \cdot \mathsf{activedegree}_\nu\right] \\ &\leqslant \Pr_g\left[|\mathsf{score}^{(1)}(g) - \mathbf{E}_g[\mathsf{score}^{(1)}(g)]| > \frac{1}{64} \cdot \mathsf{activedegree}_\nu\right] \\ &\leqslant 2 \cdot \exp\left(\frac{-2 \cdot \mathsf{activedegree}_\nu^2}{(64)^2 \sum_{v \in V\setminus S^\star} c_v^2}\right) \\ &\leqslant 2 \cdot \exp\left(\frac{-2 \cdot \mathsf{activedegree}_\nu^2}{(64)^2 \cdot (\max_v c_v) \cdot \sum_{v \in V\setminus S^\star} c_v}\right) \\ &\leqslant 2 \cdot \exp\left(\frac{-2 \cdot \mathsf{activedegree}_\nu^2}{(64)^2 \cdot (\max_v c_v) \cdot \mathsf{activedegree}_\nu}\right) \\ &\leqslant 2 \cdot \exp\left(\frac{-2 \cdot \mathsf{activedegree}_\nu}{(64)^2 \cdot (\frac{\mathsf{activedegree}_\nu}{100tk})}\right) \leqslant 2 \cdot \exp\left(\frac{-200tk}{(64)^2}\right) \leqslant 2 \cdot \exp\left(\frac{-tk}{100}\right).\end{aligned}$$

■

[6]In this case we could have simply used a Hoeffding-like inequality, but later when we handle larger-width constraints we will truly use the added generality of McDiarmid's inequality.

For a high variance instance ℓ, let $\nu_1^\ell, \ldots, \nu_t^\ell$ be the sequence of t nodes with $\mathcal{I}_\nu = \ell$ which lie on the path from the root to $\nu^\star$. Set $\mathsf{finalwt}_\ell = \mathsf{activedegree}_{\rho^\star}(\mathcal{W}_\ell)$ (in words: this is the active degree left over in instance ℓ after the restriction $\rho^\star$).

Lemma 3.6 *For every high variance instance $\ell \in [k]$ and for each $i \leqslant [t/2]$,*

$$\mathsf{activedegree}_{\nu_i^\ell} \geqslant \gamma \cdot (1-\gamma)^{-t/2} \cdot \mathsf{finalwt}_\ell \geqslant 1600tk \cdot \mathsf{finalwt}_\ell.$$

Proof: Fix a high variance instance $\ell \in [k]$. Note that $b_i = \mathsf{activedegree}_{\rho_{\nu_i^\ell}}(\mathcal{W}_\ell)$ decreases as i increases. The main observation is that

1. $b_{i+1} \leqslant (1-\gamma) \cdot b_i$.
2. $\mathsf{activedegree}_{\nu_i^\ell} \geqslant \gamma b_i$.

Thus for all ν_i^ℓ with $i \in \{1, \ldots, t/2\}$, we have $\mathsf{activedegree}_{\nu_i^\ell} \geqslant \gamma \cdot (1-\gamma)^{-t/2} \cdot \mathsf{finalwt}_\ell$ and also the choice of parameters implies for those ν_i^ℓ $\mathsf{activedegree}_{\nu_i^\ell}$ is at least $1600tk \cdot \mathsf{finalwt}_\ell$. ■

3.4.3 Putting everything together

We now show that when ν is taken to equal $\nu^\star$ in Step 5, then with high probability over the choice of g in Step 5(a) there is a setting of h in Step 5(b) such that $\forall \ell \in [k], \mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{2} - \varepsilon) \cdot c_\ell$.

Theorem 3.7 *Suppose the algorithm* Sim-Max2AND *is given as inputs $\varepsilon > 0$, k simultaneous weighted* Max-2-AND *instances $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on n variables, and target objective value $c_1, \ldots, c_k$ with the guarantee that there exists an assignment $f^\star$ such that for each $\ell \in [k]$, we have $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$. Then, the algorithm runs in $2^{O(k^4/\varepsilon^2 \log(k/\varepsilon))} \cdot \mathrm{poly}(n)$ time, and with probability at least 0.9, outputs an assignment f such that for each $\ell \in [k]$, we have, $\mathsf{val}(f, \mathcal{W}_\ell) \geqslant \left(\frac{1}{2} - \varepsilon\right) \cdot c_\ell$.*

Proof: Consider the case when ν is taken to equal $\nu^\star$ in Step 5. By Lemma 3.4, with probability at least $1 - k\delta_0$ over the choice random choices of $g_{\nu^\star}$, we have that for *every* low variance instance $\ell \in [k]$, $\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot c_\ell$. By Lemma 3.5 and a union bound, with probability at least $1 - \frac{t}{2} \cdot k \cdot 2e^{-tk/100} \geqslant 1 - \delta_0$ over the choice of $g_{\nu^\star}$, for every high variance instance ℓ and for every typeC node ν_i^ℓ, $i \in [t/2]$, we have that $g_{\nu^\star}$ is Cgood for ν_i^ℓ. Thus with probability at least $1 - (k+1)\delta_0$, both these events occur. Henceforth we assume that both these events occur in Step 5(a) of the algorithm.

Our next goal is to show that there exists a partial assignment $h : S^\star \to \{0,1\}$ such that:

1. For every instance $\ell \in [k]$, $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(1 - \frac{\varepsilon}{2}\right) \cdot \mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell)$.
2. Moreover, for every high variance instance $\ell \in [k]$, $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(1 - \frac{\varepsilon}{2}\right) \cdot \mathsf{finalwt}_\ell$.

Before giving a proof of the existence of such an h, we show that this completes the proof of the theorem. We claim that when the partial assignment h guaranteed above is considered in the Step 5(b) in the algorithm, we obtain an assignment with the required approximation guarantees.

For every low variance instance $\ell \in [k]$, since we started with $\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{2} - \frac{\varepsilon}{2}) \cdot c_\ell$, property 1 above implies that every low variance instance $\mathsf{val}(h \cup g_{\nu^\star}) \geqslant (\frac{1}{2} - \varepsilon) \cdot c_\ell$. For every high variance instance $\ell \in [k]$, since $h^\star = f^\star|_S$,

$$\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{val}(f^\star, \mathcal{W}_\ell) - \mathsf{activedegree}_{\rho^\star}(\mathcal{W}_\ell) \geqslant c_\ell - \mathsf{finalwt}_\ell.$$

Combining this with properties 1 and 2 above, we get,

$$\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \max\{c_\ell - \mathsf{finalwt}_\ell, \mathsf{finalwt}_\ell\} \geqslant 1/2 \cdot (1 - \varepsilon/2) \cdot c_\ell.$$

Thus, for all instances $\ell \in [k]$, we get $\mathsf{val}(h \cup g_{\nu^\star}) \geqslant (1/2 - \varepsilon) \cdot c_\ell$.

Now, it remains to show the existence of such an h by giving a procedure for constructing h by perturbing $h^\star$ (Note that this procedure is only part of the analysis). For nodes ν, ν' in the tree, let us write $\nu \prec \nu'$ if ν is an ancestor of ν', and we also say that ν' is "deeper" than ν.

Constructing h:

1. Initialize $H \subseteq [k]$ to be the set of high variance instances.
2. Let $N_0 = \{\nu_i^\ell \mid \ell \in H, i \in [t/2]\}$. Note that N is a chain in the tree (since all the elements of N are ancestors of $\nu^\star$). Since every $\nu \in N$ is an ancestor of $\nu^\star$, we have $h_{\rho_\nu} = h^\star|_{S_{\rho_\nu}}$.
3. Initialize $D = \varnothing$, $N = N_0$, $h = h^\star$.
4. During the procedure, we will be changing the assignment h, and removing elements from N. We will always maintain the following two invariants:
   - $|N| > \frac{t}{4}$.
   - For every $\nu \in N$, $h|_{S_{\rho_\nu}} = h^\star|_{S_{\rho_\nu}}$.
5. While $|D| \neq |H|$ do:
   (a) Let
   $$B = \left\{ v \in V \mid \exists \ell \in [k] \text{ with } \sum_{C \in \mathcal{C}, C \ni v} \mathcal{W}_\ell(C) \cdot C(h \cup g_{\nu^\star}) \geqslant \frac{\varepsilon}{4k} \mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \right\}.$$
   Note that $|B| \leqslant \frac{8k^2}{\varepsilon} < \frac{t}{8}$.
   (b) Let $\nu \in N$ be the deepest element of N for which: $\{\mathcal{T}_\nu^1, \mathcal{T}_\nu^2\} \cap B = \varnothing$.
   Such a ν exists because:
   - $|N| > \frac{t}{4} > |B|$, and
   - there are at most $|B|$ nodes ν for which $\{\mathcal{T}_\nu^1, \mathcal{T}_\nu^2\} \cap B \neq \varnothing$ (since $\{\mathcal{T}_\nu^1, \mathcal{T}_\nu^2\}$ are all disjoint for distinct ν).

   (c) Let $\ell \in H, i \in [t/2]$ be such that $\nu = \nu_i^\ell$. Let $x = \mathcal{T}_\nu^1$ and $y = \mathcal{T}_\nu^2$. Let $\rho = \rho_\nu$. We will now see a way of modifying the values of $h(x)$ and $h(y)$ to guarantee that $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{finalwt}_\ell$. The procedure depends on whether ν is typeA, typeB, or typeC.

i. If ν is typeA, then we know that $\mathsf{backward}_\nu \geqslant \frac{1}{2} \cdot \mathsf{activedegree}_\nu \geqslant 2 \cdot \mathsf{finalwt}_\ell$. The second invariant tells us that $\rho = h^\star|_{S_\rho} = h|_{S_\rho}$. Thus we have:

$$\begin{aligned}\mathsf{backward}_\nu &= \sum_{C \in \mathcal{C}_\nu^{\mathsf{backward}}} \mathcal{W}_\ell(C) \\ &= \sum_{C \subseteq S_\rho \cup \{x\}, C \ni x, C \in \mathsf{Active}(\rho)} \mathcal{W}_\ell(C) \\ &= \sum_{C \subseteq S_\rho \cup \{x\}, C \ni x, C \in \mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C).\end{aligned}$$

This implies that we can choose a setting of $h(x) \in \{0,1\}$ such that the total sum of weights of those constraints containing x which are satisfied by h is:

$$\begin{aligned}\sum_{C \subseteq S_\rho \cup \{x\}, C \ni x, C \in \mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C) C(h) &\geqslant \frac{1}{2} \sum_{C \subseteq S_\rho \cup \{x\}, C \ni x, C \in \mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C) \\ &= \frac{1}{2} \cdot \mathsf{backward}_\nu \\ &\geqslant \frac{1}{4} \cdot \mathsf{activedegree}_\nu \\ &\geqslant \mathsf{finalwt}_\ell, \quad \text{(by Lemma 3.6)}\end{aligned}$$

where the $\frac{1}{2}$ in the first inequality is because the variable can appear as a *positive* literal or a *negative* literal in those $\mathsf{backward}$ constraints. In particular, after making this change, we have $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{finalwt}_\ell$.

ii. If ν is typeB, then we know that some constraint C containing x and y has $\mathcal{W}_\ell(C) \geqslant \frac{1}{1600tk} \cdot \mathsf{activedegree}_\nu \geqslant \mathsf{finalwt}_\ell$. Thus we may choose settings for $h(x), h(y) \in \{0,1\}$ such that $C(h) = 1$. Thus, after making this assignment to $h(x)$ and $h(y)$, we have $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant \mathsf{finalwt}_\ell$.

iii. If ν is typeC, since $g_{\nu^\star}$ is Cgood for ν, we can choose a setting of $h(x)$ so that the total weight of satisfied constraints in $\mathcal{W}_\ell$ between x and $V \backslash S^\star$ is at least $\frac{1}{64} \cdot \mathsf{activedegree}_\nu \geqslant \mathsf{finalwt}_\ell$. After this change, we have $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{finalwt}_\ell$.

In all the above 3 cases, we only changed the value of h at the variables x, y. Since $\{x, y\} \cap B = \varnothing$, we have that for every $j \in [k]$, the new value $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_j)$ is at least $\left(1 - \frac{\varepsilon}{2k}\right)$ times the old value $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_j)$.

(d) Set $D = D \cup \{\ell\}$.

(e) Set $N = \{\nu_i^\ell \mid \ell \in H \backslash D, i \leqslant [t/2], \nu_i^\ell \prec \nu\}$.

Observe that $|N|$ decreases in size by at most $\frac{t}{2} + |B|$. Thus, if $D \neq H$, we have

$$\begin{aligned}|N| &\geqslant |N_0| - |D| \cdot \frac{t}{2} - |D||B| \\ &= |H| \cdot \frac{t}{2} - |D| \cdot \frac{t}{2} - |D||B| \\ &\geqslant \frac{t}{2} - k|B| > \frac{t}{4}\end{aligned}$$

Also observe that we only changed the values of h at the variables $\mathcal{T}_\nu^1$ and $\mathcal{T}_\nu^2$. Thus for all $\nu' \preceq \nu$, we still have the property that $h|_{S_{\rho_{\nu'}}} = h^\star|_{S_{\rho_{\nu'}}}$.

For each high variance instance $\ell \in [k]$, in the iteration where ℓ gets added to the set D, the procedure ensures that at the end of the iteration $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{finalwt}_\ell$.

Moreover, at each step we reduced the value of each $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell)$ by at most $\frac{\varepsilon}{2k}$ fraction of its previous value. Thus, at the end of the procedure, for every $\ell \in [k]$, the value has decreased at most by a multiplicative factor of $\left(1 - \frac{\varepsilon}{2k}\right)^k \geqslant \left(1 - \frac{\varepsilon}{2}\right)$. Thus, for every $\ell \in [k]$, we get $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(1 - \frac{\varepsilon}{2}\right) \cdot \mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell)$, and for every high variance instance $\ell \in [k]$, we have $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(1 - \frac{\varepsilon}{2}\right) \cdot \mathsf{finalwt}_\ell$. This proves the two properties of h that we set out to prove.

Running time : Running time of the algorithm is $2^{O(kt)} \cdot \mathrm{poly}(n)$ which is $2^{O(k^4/\varepsilon^2 \log(k/\varepsilon^2))} \cdot \mathrm{poly}(n)$. ■

4 Simultaneous Max-w-CSP$_q$

In this section, we give our simultaneous approximation algorithm for Max-w-CSP$_q$, and thus prove Theorem 1.2.

4.1 Reduction to simple constraints

For the problem Max-w-CSP$_q$, $\mathcal{C}$ is the set of all possible q-ary constraints on V with arity at most w, *i.e.*, each constraint is of the form $C_f : [q]^T \to \{0, 1\}$ depending only on the values of variables in an ordered tuple $T \subseteq V$ with $|T| \leqslant w$. As a first step (mainly to simplify notation), we give a simple approximation preserving reduction which replaces $\mathcal{C}$ with a smaller set of constraints. We will then present our main algorithm

Define a *w-term* to be a contraint C on exactly w variables which has exactly 1 satisfying assignment in $[q]^w$, *e.g.* $(x_1 = 1) \wedge (x_2 = 7) \wedge \ldots \wedge (x_w = q - 3)$. An instance of the Max-w-ConjSAT$_q$ problem is one where the set of constraints $\mathcal{C}$ is the set of all w-terms. We now use the following lemma from [Tre98] that allows us to reduce a Max-w-CSP$_q$ instance to a Max-w-ConjSAT$_q$ instance.

Lemma 4.1 ([Tre98]) *Given an instance $\mathcal{W}_1$ of* Max-w-CSP$_q$*, we can find a instance $\mathcal{W}_2$ of* Max-w-Conj-Sat$_q$ *on the same set of variables, and a constant $\beta > 0$ such that for every assignment f, $\mathsf{val}(f, \mathcal{W}_2) = \beta \cdot \mathsf{val}(f, \mathcal{W}_1)$.*

Proof: Given an instance $\mathcal{W}_1$ of Max-w-CSP$_q$, consider a constraint $C \in \mathcal{C}$ with weight $\mathcal{W}_1(C)$. We can assume without loss of generality that the arity of C is exactly w, and it depends on variables $x_1, \ldots, x_w$. For each assignment in $[q]^k$ that satisfies C, we create a w-Conj-Sat$_q$ clause that is satisfied only for that assignment, and give it weight $\mathcal{W}_1(C)$. *e.g.* If C was satisfied by $x_1 = \ldots = x_w = 2$, we create the clause $(x_1 = 2) \wedge \ldots \wedge (x_w = 2)$ with weight $\mathcal{W}_1(C)$. It is easy to see that for every assignment to $x_1, \ldots, x_n$, the weight of constraints satisfied in the new instance is the same as the weight of the constraints satisfied in the Max-w-Conj-Sat$_q$ instance created. Define β to be the sum of weights of all the constraints in the new instance, then $\mathcal{W}_2$ is obtained by multiplying the weight of all the constraints in the new instance by $1/\beta$ (to make sure they sum up to 1). ■

Note that the scaling factor β in the lemma above is immaterial since we will give an algorithm with Pareto approximation guarantee.

We say $(v, i) \in C$ if $v \in C$ and $v = i$ is in the satisfying assignment of C_f. By abuse of notation, we say for a set of variables T, $T \subseteq C$ if for all $v \in T$, there exists $i \in [q]$, such that $(v, i) \in C$.

4.2 Random Assignments

In this section, we state and prove a lemma that gives a sufficient condition for the value of a MAX-w-CONJSAT$_q$ to be highly concentrated under independent random assignments to the variables. Let $\mathsf{Dist}(q)$ denote the set of all probability distributions on the set $[q]$. For a distribution $p \in \mathsf{Dist}(q)$ and $i \in q$, we use p_i to denote the probability i in the distribution p. Let ρ be a partial assignment. Let $p : V\backslash S_\rho \to \mathsf{Dist}(q)$ be such that $p(v)_i \geqslant \frac{1}{qw}$ for all $v \in V\backslash S_\rho$ and all $i \in [q]$. Let $g : V\backslash S_\rho \to [q]$ be a random assignment obtained by sampling $g(v)$ for each v independently according to the distribution $p(v)$.

Define the random variable

$$Y \stackrel{\text{def}}{=} \mathsf{val}(\rho \cup g, \mathcal{W}) - \mathsf{val}(\rho, \mathcal{W}) = \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho \cup g).$$

The random variable Y measures the contribution of active constraints to $\mathsf{val}(\rho \cup g, \mathcal{W})$. Note that the two quantities $\mathbf{E}[Y]$ and $\mathbf{Var}[Y]$ can be computed efficiently given p. We denote these by $\mathsf{TrueMean}_\rho(p, \mathcal{W})$ and $\mathsf{TrueVar}_\rho(p, \mathcal{W})$. The following lemma is a generalization of Lemma 3.1.

Lemma 4.2 *Let p, g, Y be as above.*

1. *If $\mathsf{TrueVar}_\rho(p, \mathcal{W}) < \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_\rho(p, \mathcal{W})^2$ then $\Pr[Y < (1 - \varepsilon_0)\mathbf{E}[Y]] < \delta_0$.*
2. *If $\mathsf{TrueVar}_\rho(p, \mathcal{W}) \geqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_\rho(p, \mathcal{W})^2$, then there exists $v \in V\backslash S_\rho$ such that*

$$\mathsf{activedegree}_\rho(v, \mathcal{W}) \geqslant \frac{\varepsilon_0^2 \delta_0}{w^2 (qw)^w} \cdot \mathsf{activedegree}_\rho(\mathcal{W}).$$

Proof: Item 1 of the lemma follows immediately from Chebyshev's inequality. We now prove Item 2. First note that for every active constraint C given ρ, $\mathbf{E}[C(\rho \cup g)] \geqslant \frac{1}{(qw)^w}$ (this follows from our hypothesis that $p(v)_i \geqslant \frac{1}{qw}$ for each $v \in V\backslash S_\rho$ and each $i \in [q]$).

We first bound $\mathsf{TrueMean}_\rho(p, \mathcal{W})$ and $\mathsf{TrueVar}_\rho(p, \mathcal{W})$ in terms of the weights of active constraints:

$$\begin{aligned}
\mathsf{TrueMean}_\rho(p, \mathcal{W}) = \mathbf{E}[Y] &= \mathbf{E}\left[\sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) \cdot C(\rho \cup g)\right] \\
&= \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) \cdot \mathbf{E}[C(\rho \cup g)] \geqslant \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) \cdot \frac{1}{(qw)^w} \\
&= \frac{1}{(qw)^w} \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C)
\end{aligned}$$

$$\begin{aligned}
\mathsf{TrueVar}_\rho(p, \mathcal{W}) = \mathbf{Var}[Y] &= \mathbf{Var}\left[\sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) \cdot C(\rho \cup g)\right] \\
&= \sum_{C_1, C_2 \in \mathsf{Active}(\rho)} \mathcal{W}(C_1)\mathcal{W}(C_2) \cdot (\mathbf{E}[C_1(\rho \cup g) C_2(\rho \cup g)] - \mathbf{E}[C_1(\rho \cup g)]\mathbf{E}[C_2(\rho \cup g)])
\end{aligned}$$

$$
\begin{aligned}
&\leqslant \sum_{C_1 \sim_\rho C_2} \mathcal{W}(C_1)\mathcal{W}(C_2) \cdot \mathbf{E}[C_1(\rho \cup g)] \\
&= \sum_{C_1 \in \mathsf{Active}(\rho)} \mathcal{W}(C_1)\mathbf{E}[C_1(\rho \cup g)] \cdot \sum_{C_2 \sim_\rho C_1} \mathcal{W}(C_2) \\
&\leqslant \sum_{C_1 \in \mathsf{Active}(\rho)} \mathcal{W}(C_1)\mathbf{E}[C_1(\rho \cup g)] \cdot \max_{C \in \mathsf{Active}(\rho)} \sum_{C_2 \sim_\rho C} \mathcal{W}(C_2) \\
&= \mathsf{TrueMean}_\rho(p, \mathcal{W}) \cdot \max_{C \in \mathsf{Active}(\rho)} \sum_{C_2 \sim_\rho C} \mathcal{W}(C_2).
\end{aligned}
$$

Hence, if the condition in case 2 is true then it follows that,

$$
\max_{C \in \mathsf{Active}(\rho)} \sum_{C_2 \sim_\rho C} \mathcal{W}(C_2) \geqslant \frac{\mathsf{TrueVar}_\rho(p, \mathcal{W})}{\mathsf{TrueMean}_\rho(p, \mathcal{W})} \geqslant \frac{\delta_0 \varepsilon_0^2}{(qw)^w} \cdot \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C).
$$

We now relate these quantities to active degrees.

$$
\begin{aligned}
\mathsf{activedegree}_\rho(\mathcal{W}) &= \sum_{v \in V \setminus S_\rho} \mathsf{activedegree}_\rho(v, \mathcal{W}) = \sum_{v \in V \setminus S_\rho} \sum_{C \in \mathsf{Active}(\rho), C \ni v} \mathcal{W}_\ell(C) \\
&= \sum_{C \in \mathsf{Active}(\rho)} \sum_{v \in C, v \in V \setminus S_\rho} \mathcal{W}_\ell(C) \leqslant \sum_{C \in \mathsf{Active}(\rho)} w \cdot \mathcal{W}_\ell(C) \\
&= w \sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}_\ell(C)
\end{aligned}
$$

This means that there is an active constraint C, such that

$$
\sum_{C_2 \sim_\rho C} \mathcal{W}(C_2) \geqslant \frac{\delta_0 \varepsilon_0^2}{(qw)^w} \cdot \frac{1}{w} \mathsf{activedegree}_\rho(\mathcal{W})
$$

Since C is an active constraint and $|C \cap V \setminus S_\rho| \leqslant w$, there is some variable $v \in C \cap V \setminus S_\rho$, such that

$$
\mathsf{activedegree}_\rho(v, \mathcal{W}) = \sum_{C_2 \in \mathsf{Active}(\rho),\ C_2 \ni v} \mathcal{W}(C_2) \geqslant \frac{1}{w} \sum_{C_2 \sim_\rho C} \mathcal{W}(C_2) \geqslant \frac{\varepsilon_0^2 \delta_0}{w^2 (qw)^w} \cdot \mathsf{activedegree}_\rho(\mathcal{W}).
$$

as required. ■

4.3 LP relaxations

Our algorithm will use the Linear Programming relaxation for MAX-w-CONJSAT$_q$ from the work of Trevisan [Tre98] (actually, a simple generalization to q-ary alphabets). The first LP, $\mathsf{ConjSAT\text{-}LP}_1(\rho)$, described in Figure 4, describes the set of all feasible solutions for the relaxation, consistent with the partial assignment ρ. Given a set of target values $(c_\ell)_{\ell \in [k]}$, the second LP, $\mathsf{ConjSAT\text{-}LP}_2(\rho)$ describes the set of feasible solutions to $\mathsf{ConjSAT\text{-}LP}_1(\rho)$ that achieve the required objective values.

For $\vec{t}$, $\vec{z}$ satisfying linear constraints $\mathsf{ConjSAT\text{-}LP}_1(\rho)$, let $\mathsf{smooth}(\vec{t})$ denote the map $p : V \setminus S_\rho \to \mathsf{Dist}(q)$ with $p(v)_i = \frac{w-1}{qw} + \frac{t_{v,i}}{w}$. The following theorem from [Tre98] provides an algorithm to round this feasible solution to obtain a good integral assignment.

Lemma 4.3 *Let ρ be a partial assignment.*

1. **Relaxation:** *For every $g_0 : V \backslash S_\rho \to [q]$, there exist $\vec{t}$, $\vec{z}$ satisfying* $\mathsf{ConjSAT\text{-}LP}_1(\rho)$ *such that for every* Max-w-ConjSAT$_q$ *instance $\mathcal{W}$:*

$$\sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C = \mathsf{val}(g_0 \cup \rho, W).$$

2. **Rounding:** *Suppose $\vec{t}, \vec{z}$ satisfy* $\mathsf{ConjSAT\text{-}LP}_1(\rho)$. *Then for every* Max-$w$-ConjSAT$_q$ *instance $\mathcal{W}$:*

$$\mathsf{val}(\rho, \mathcal{W}) + \mathsf{TrueMean}_\rho(\mathsf{smooth}(\vec{t}), \mathcal{W}) \geqslant \frac{1}{q^{w-1}} \cdot \sum_{C \in \mathcal{C}} z_C \mathcal{W}(C).$$

Proof: We begin with the first part. For $v \in S_\rho$, $i \in [q]$, define $t_{v,i} = 1$ if $\rho(v) = i$, and define $t_{v,i} = 0$ otherwise. For $v \in V \backslash S_\rho, i \in [q]$, define $t_{v,i} = 1$ if $g_0(v) = i$, and define $t_{v,i} = 0$ otherwise. For $C \in \mathcal{C}$, define $z_C = 1$ if $C(g_0 \cup \rho) = 1$, and define $z_C = 0$ otherwise. It is easy to see that these $\vec{t}, \vec{z}$ satisfies $\mathsf{ConjSAT\text{-}LP}_1(\rho)$, and that for every instance $\mathcal{W}$:

$$\sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C = \mathsf{val}(g_0 \cup \rho, W).$$

Now we consider the second part. Let $\mathcal{W}$ be any instance of Max-w-ConjSAT$_q$. Let $p = \mathsf{smooth}(t)$. Let $g : V \backslash S_\rho \to [q]$ be sampled as follows: independently for each $v \in V \backslash S_\rho$, $g(v)$ is sampled from the distribution $p(v)$. We need to show that:

$$\sum_{C \notin \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho) + \mathbf{E}\left[\sum_{C \in \mathsf{Active}(\rho)} \mathcal{W}(C) C(\rho \cup g) \right] \geqslant \frac{1}{q^{w-1}} \cdot \sum_{C \in \mathcal{C}} z_C \mathcal{W}(C).$$

It is easy to check that for $C \notin \mathsf{Active}(\rho)$, $z_C > 0$ only if $C(\rho) = 1$, and thus $\sum_{C \notin \mathsf{Active}(\rho)} C(\rho) \mathcal{W}(C) \geqslant \sum_{C \notin \mathsf{Active}(\rho)} z_C \mathcal{W}(C)$. For $C \in \mathsf{Active}(\rho)$, we have the following claim:

Claim 4.4 *For $C \in \mathsf{Active}(\rho)$, $\mathbf{E}[C(\rho \cup g)] \geqslant \frac{z_C}{q^{w-1}}$.*

Proof: Suppose there are exactly h variables in C which are not in S_ρ. Let these variables be $(v_i)_{i=1}^h$. Let $(v_i, a_i)_{i=1}^h$ be the assignment to these variables that makes C satisfied.

$$\begin{aligned}
\mathbf{E}[C(\rho \cup g)] = \Pr[C \text{ is satisfied by } \rho \cup g] &\geqslant \prod_{i=1}^{h} \left(\frac{w-1}{qw} + \frac{t_{v_i,a_i}}{w} \right) \\
&\geqslant \prod_{i=1}^{h} \left(\frac{w-1}{qw} + \frac{z_C}{w} \right) = \left(\frac{w-1}{qw} + \frac{z_C}{w} \right)^h \\
&= \left(\frac{w-1}{qw} + \frac{z_C}{w} \right)^w \geqslant \frac{z_C}{q^{w-1}}
\end{aligned}$$

Here the last inequality follows form the observation that the minimum of the function $\frac{\left(\frac{w-1}{qw} + \frac{z}{w}\right)^w}{z}$ as z varies in $[0, 1]$, is attained at $z = 1/q$. ■

This completes the proof of the Lemma.

■

$$
\begin{array}{rcll}
z_C & \leqslant & t_{v,i} & \forall C \in \mathcal{C}, \forall (v,i) \in C \\
1 \geqslant t_{v,i} & \geqslant & 0 & \forall v \in V \backslash S_\rho, i \in [q] \\
\sum_{i=1}^{q} t_{v,i} & = & 1 & \forall v \in V \\
t_{v,i} & = & 1 & \forall v \in S_\rho \text{ and } i \in [q], \\
 & & & \text{such that } h_\rho(v) = i
\end{array}
$$

Figure 4: Linear inequalities $\mathsf{ConjSAT\text{-}LP}_1(\rho)$

$$
\begin{array}{c}
\sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) \cdot z_C \geqslant c_\ell \quad \forall \ell \in [k] \\
\vec{t}, \vec{z} \text{ satisfy } \mathsf{ConjSAT\text{-}LP}_1(\rho).
\end{array}
$$

Figure 5: Linear inequalities $\mathsf{ConjSAT\text{-}LP}_2(\rho)$

4.4 The Algorithm

We now give our Pareto approximation algorithm for MAX-w-CSP$_q$ in Figure 7 (which uses the procedure from Figure 6).

Input: A tree node ν and an instance $\mathcal{W}_\ell$.
Output: A tuple of variables of size at most w.

1. Let $v_1 \in V \backslash S_{\rho_\nu}$ be a variable which maximizes the value of $\mathsf{activedegree}_{\rho_\nu}(v_1, \mathcal{W}_\ell)$. Set $D \leftarrow \{v_1\}$.
2. While $|D| \leqslant w$, do the following
   (a) If there is a variable v in $V \backslash S_{\rho_\nu}$ such that
   $$\mathsf{activedegree}_{\rho_\nu}(D \cup v, \mathcal{W}_\ell) \geqslant \frac{\mathsf{activedegree}_{\rho_\nu}(D, \mathcal{W}_\ell)}{(4qwtk)^w},$$
   set $D \leftarrow D \cup v$.
   (b) Otherwise, go to Step 3.
3. Return D as a tuple (in arbitrary order, with v_1 as the first element).

Figure 6: TUPLESELECTION for MAX-w-CONJSAT$_q$

Input: k instances of MAX-w-CONJSAT$_q$ $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on the variable set V, $\varepsilon > 0$ and and target objective values $c_1, \ldots, c_k$.
Output: An assignment to V
Parameters: $\delta_0 = \frac{1}{10(k+1)}$, $\varepsilon_0 = \varepsilon$, $\gamma = \frac{\varepsilon_0^2 \delta_0}{w^2 (qw)^w}$, $t = \left\lceil \frac{20 w^2 k^2}{\gamma} \cdot \log\left(\frac{10k}{\gamma}\right) \right\rceil$

1. Initialize tree T to be an empty q^w-ary tree (i.e., just 1 root node and each node has at most q^w children).
2. We will associate with each node ν of the tree:
   (a) A partial assignment ρ_ν.
   (b) A special set of variables $\mathcal{T}_\nu \subseteq V \backslash S_{\rho_\nu}$.
   (c) A special instance $\mathcal{I}_\nu \in [k]$.
   (d) A collection of integers $\mathsf{count}_{\nu,1}, \ldots, \mathsf{count}_{\nu,k}$.
   (e) A trit representing whether the node ν is living, exhausted or dead.
3. Initialize the root node ν_0 to (1) $\rho_{\nu_0} = (\varnothing, \varnothing)$, (2) have all $\mathsf{count}_{\nu_0,\ell} = 0$, (3) living.
4. While there is a living leaf ν of T, do the following:
   (a) Check if the LP $\mathsf{ConjSAT\text{-}LP}_2(\rho_\nu)$ has a feasible solution.
      i. If $\vec{t}, \vec{z}$ is a feasible solution, then define $p_\nu : V \backslash S_{\rho_\nu} \to \mathsf{Dist}(q)$ by $p = \mathsf{smooth}(\vec{t})$.
      ii. If not, then declare ν to be dead and return to Step 4.
   (b) For each $\ell \in [k]$, compute $\mathsf{TrueVar}_{\rho_\nu}(p, \mathcal{W}_\ell), \mathsf{TrueMean}_{\rho_\nu}(p, \mathcal{W}_\ell)$.
   (c) If $\mathsf{TrueVar}_{\rho_\nu}(p, \mathcal{W}_\ell) \geqslant \delta_0 \varepsilon_0^2 \mathsf{TrueMean}_{\rho_\nu}(p, \mathcal{W}_\ell)^2$, then set $\mathsf{flag}_\ell = \textsc{True}$, else set $\mathsf{flag}_\ell = \textsc{False}$.
   (d) Choose the smallest $\ell \in [k]$, such that $\mathsf{count}_\ell < t$ AND $\mathsf{flag}_\ell = \textsc{True}$ (if any):
      i. Set $\mathcal{T}_\nu \leftarrow \textsc{TupleSelection}(\nu, \mathcal{W}_\ell)$. Set $\mathcal{I}_\nu = \ell$.
      ii. Create $q^{w'}$ children of ν, with labels νb for each $b \in [q]^{w'}$ and define
         – $\rho_{\nu b} = (S_{\rho_\nu} \cup \mathcal{T}_\nu, h^b)$, where h^b extends h_{ρ_ν} by $h^b(\mathcal{T}_\nu^i) = b(i)$.
         – For each $\ell' \in [k]$ with $\ell' \neq \ell$, initialize $\mathsf{count}_{\nu b,\ell'} = \mathsf{count}_{\nu,\ell'}$. Initialize $\mathsf{count}_{\nu b,\ell} = \mathsf{count}_{\nu,\ell} + 1$.
         – Set νb to be living.
   (e) If no such ℓ exists, declare ν to be exhausted.
5. Now every leaf of T is either exhausted or dead. For each exhausted leaf ν of T:
   (a) Sample $g_\nu : V \backslash S_{\rho_\nu} \to [q]$ by independently sampling $g_\nu(v)$ from the distribution $p_\nu(v)$.
   (b) For every assignment $h : S_{\rho_\nu} \to [q]$, compute $\mathsf{out}_{h,g_\nu} \leftarrow \min_{\ell \in [k]} \frac{\mathsf{val}(h \cup g_\nu, \mathcal{W}_\ell)}{c_\ell}$. If $c_\ell = 0$ for some $\ell \in [k]$, we interpret $\frac{\mathsf{val}(h \cup g_\nu, \mathcal{W}_l)}{c_\ell}$ as $+\infty$.
6. Output the largest out_{h,g_ν} seen, and the assignment $h \cup g_\nu$ that produced it.

Figure 7: Algorithm SIM-MAXCONJSAT$_{26}$ for approximating weighted simultaneous MAX-w-CONJSAT$_q$

4.5 Analysis

Notice that the depth of the tree T is at most kt, and that for every ν, we have that $|S_{\rho_\nu}| \leqslant wkt$. This implies that the running time is at most $q^{O(wkt)} \cdot \mathrm{poly}(n)$.

Let $f^\star : V \to [q]$ be an assignment such that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for each $\ell \in [k]$. Let $\nu^\star$ be the the unique leaf of the tree T for which $\rho_{\nu^\star}$ is consistent with $f^\star$. (This $\nu^\star$ can be found as follows: start with ν equal to the root. Set ν to equal the unique child of ν for which ρ_ν is consistent with $f^\star$, and repeat until ν becomes a leaf. This leaf is $\nu^\star$). Observe that since $f^\star$ is an assignment such that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for every $\ell \in [k]$, by picking $g_0 = f^\star|_{V \setminus S^\star}$ in part 1 of Lemma 4.3, we know that $\mathsf{ConjSAT\text{-}LP}_2(\rho^\star)$ is feasible, and hence $\nu^\star$ must be an exhausted leaf (and not dead).

Define $\rho^\star = \rho_{\nu^\star}$, $S^\star = S_{\rho^\star}$, $h^\star = h_{\rho^\star}$ and $p^\star = p_{\nu^\star}$

At the completion of Step 4, if $\ell \in [k]$ satisfies $\mathsf{count}_{\nu^\star,\ell} = t$, we call instance ℓ a *high variance* instance. Otherwise we call instance ℓ a *low variance* instance.

4.5.1 Low Variance Instances

First we show that for the leaf ν^* in Step 5, combining the partial assignment $h^\star$ with a random assignment $g_{\nu^\star}$ in Step 5a is good for any low variance instances with high probability.

Lemma 4.5 *Let $\ell \in [k]$ be any low variance instance. For the leaf node $\nu^\star$, let $g_{\nu^\star}$ be the random assignment sampled in Step 5a of* Sim-MaxConjSAT. *Then with probability at least $1 - \delta_0$, the assignment $f = h^\star \cup g_{\nu^\star}$ satisfies:*

$$\Pr_{g_{\nu^\star}}\left[\mathsf{val}(f, \mathcal{W}_\ell) \geqslant (1/q^{w-1} - \varepsilon/2) \cdot c_\ell\right] \geqslant 1 - \delta_0.$$

Proof: For every low variance instance ℓ, we have that $\mathsf{TrueVar}_{\rho_{\nu^\star}}(p^\star, \mathcal{W}_\ell) < \delta_0 \varepsilon_0^2 \cdot \mathsf{TrueMean}_{\rho_{\nu^\star}}(p^\star, \mathcal{W}_\ell)^2$. Define $Y \stackrel{\text{def}}{=} \mathsf{val}(\rho^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) - \mathsf{val}(\rho^\star, \mathcal{W}_\ell)$. By Lemma 4.2, we have $\Pr[Y < (1-\varepsilon_0)\mathbf{E}[Y]] < \delta_0$. Thus, with probability at least $1 - \delta_0$, we have,

$$\begin{aligned}
\mathsf{val}(f, \mathcal{W}_\ell) &\geqslant \mathsf{val}(\rho^\star, \mathcal{W}_\ell) + (1-\varepsilon_0)\mathbf{E}[Y] \\
&= \mathsf{val}(\rho^\star, \mathcal{W}_\ell) + (1-\varepsilon_0) \cdot \mathsf{TrueMean}_{\rho_{\nu^\star}}(\mathsf{smooth}(\vec{t}), \mathcal{W}_\ell) \\
&= (1-\varepsilon_0) \cdot \left(\mathsf{val}(\rho^\star, \mathcal{W}_\ell) + \mathsf{TrueMean}_{\rho_{\nu^\star}}(\mathsf{smooth}(\vec{t}), \mathcal{W}_\ell)\right) \\
&\geqslant \frac{1}{q^{w-1}} \cdot (1-\varepsilon_0) \cdot \sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) \cdot z_C \geqslant \frac{1}{q^{w-1}} \cdot (1-\varepsilon_0) \cdot c_\ell \geqslant \left(\frac{1}{q^{w-1}} - \frac{\varepsilon}{2}\right) \cdot c_\ell,
\end{aligned}$$

where we have used the second part of Lemma 4.3. ∎

Next, we will consider a small perturbation of $h^\star$ which will ensure that the algorithm performs well on high variance instances too. We will ensure that this perturbation does not affect the success on the low variance instances.

4.5.2 High variance instances

Fix a high variance instance ℓ. Let ν be an ancestor of $\nu^\star$ with $\mathcal{I}_\nu = \ell$. Let $\mathcal{T}_\nu^1$ denote the first element of the tuple $\mathcal{T}_\nu$. Define:

$$\mathsf{activedegree}_\nu \stackrel{\text{def}}{=} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu^1, \mathcal{W}_\ell).$$

$$\mathsf{activedegree}_{\mathcal{T}_\nu} \stackrel{\text{def}}{=} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu, \mathcal{W}_\ell).$$

Observation 4.6 *For any node ν, in the tree,*

$$\mathsf{activedegree}_{\mathcal{T}_\nu} \geqslant \frac{\mathsf{activedegree}_\nu}{(4qwtk)^{w\cdot(|\mathcal{T}_\nu|-1)}}.$$

Proof: For ν such that $|\mathcal{T}_\nu| = 1$, we have, by definition, $\mathsf{activedegree}_{\mathcal{T}_\nu} = \mathsf{activedegree}_\nu$ and the inequality follows. The lower bound is obvious from the TUPLE SELECTION procedure if $|\mathcal{T}_\nu| > 1$. ■

Let $\mathcal{C}_\nu$ be the set of all constraints C containing all variables in $\mathcal{T}_\nu$ which are active given ρ_ν.

We call a constraint C in $\mathcal{C}_\nu$ a backward constraint if C only involves variables from $S_{\rho_\nu} \cup \mathcal{T}_\nu$. Otherwise we call C in $\mathcal{C}_\nu$ a forward constraint. Let $\mathcal{C}_\nu^{\mathsf{backward}}$ and $\mathcal{C}_\nu^{\mathsf{forward}}$ denote the sets of these constraints. Finally, let $\mathcal{C}_\nu^{\mathsf{out}}$ denote the set of all constraints from $\mathcal{C}_\nu$ that involve at least *one* variable from $V\backslash S^\star$ and none from $S^\star\backslash S_{\rho_\nu}$.

Define backward degree and forward degree as follows:

$$\mathsf{backward}_\nu \stackrel{\text{def}}{=} \sum_{C\in\mathcal{C}_\nu^{\mathsf{backward}}} \mathcal{W}_\ell(C),$$

$$\mathsf{forward}_\nu \stackrel{\text{def}}{=} \sum_{C\in\mathcal{C}_\nu^{\mathsf{forward}}} \mathcal{W}_\ell(C).$$

Note that:

$$\mathsf{activedegree}_{\mathcal{T}_\nu} = \mathsf{backward}_\nu + \mathsf{forward}_\nu.$$

Based on the above definitions, we classify ν into one of three categories:

1. If $\mathsf{backward}_\nu \geqslant \frac{1}{2}\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}$, then we call ν typeAB.
2. Otherwise, we call ν typeC.

We have the following lemma about typeC nodes.

Lemma 4.7 *For every* typeC *node ν, we have*

1. *For every $v \in V\backslash(S_{\rho_\nu} \cup \mathcal{T}_\nu)$, $\mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu \cup \{v\}, \mathcal{W}_\ell) \leqslant \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w}$.*
2. *$\sum_{C\in\mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) \geqslant \frac{1}{4}\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}$.*

Proof: If ν is a typeC node, we must have that for every $v \in V\backslash(S_{\rho_\nu} \cup \mathcal{T}_\nu)$,

$$\mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu \cup \{v\}, \mathcal{W}_\ell) < \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w}.$$

This follows from the description of the TUPLESELECTION procedure, and the observation that $\mathsf{activedegree}_\nu(T, \mathcal{W}_\ell) = 0$ for any $T \subset V$ with $|T| > w$.

In particular, since $|S^\star| \leqslant wtk$, the total weight of constraints containing $\mathcal{T}_\nu$ and some variable in $S^\star \backslash (S_{\rho_\nu} \cup \mathcal{T}_\nu)$ is at most

$$\begin{aligned}\sum_{v \in S^\star \backslash (S_{\rho_\nu} \cup \mathcal{T}_\nu)} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu \cup \{v\}, \mathcal{W}_\ell) &\leqslant \sum_{v \in S^\star \backslash (S_{\rho_\nu} \cup \mathcal{T}_\nu)} \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w} \\ &\leqslant |S^\star \backslash (S_{\rho_\nu} \cup \mathcal{T}_\nu)| \cdot \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w} \\ &\leqslant wtk \cdot \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w} \leqslant \frac{1}{4} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu}.\end{aligned}$$

Thus, we get,

$$\begin{aligned}\sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) &= \mathsf{forward}_\nu - \left\{ \begin{smallmatrix} \text{total weight of constraints containing} \\ \mathcal{T}_\nu \text{ and some variable in } S^\star \backslash (S_{\rho_\nu} \cup \mathcal{T}_\nu) \end{smallmatrix} \right\} \\ &\geqslant \frac{1}{2} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu} - \frac{1}{4} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu} = \frac{1}{4} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu}.\end{aligned}$$

This completes the proof of second statement. ■

For a partial assignment $g : V \backslash S^\star \to [q]$, we say that g is Cgood for ν if there exists a setting of variables in $\mathcal{T}_\nu$ that satisfies at least $\frac{1}{8 \cdot (qw)^w} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu}$ weight amongst constraints in $\mathcal{C}_\nu^{\mathsf{out}}$.

The next lemma allows us to prove that that for every node ν of typeC, with high probability, the random assignment $g_{\nu^\star} : V \backslash S^\star \to [q]$, is Cgood for ν.

Lemma 4.8 *Let ν be typeC. Suppose $g : V \backslash S^\star \to [q]$ is a random assignment obtained by independently sampling $g(v)$ for each $v \in V \backslash S^\star$ from a distribution such that distribution $\Pr[g(v) = i] \geqslant \frac{1}{qw}$ for each $i \in [q]$. Then:*

$$\Pr_g[g \text{ is } \mathsf{Cgood} \text{ for } \nu] \geqslant 1 - 2 \cdot e^{-tk/8qw}.$$

Proof: Let $\ell = \mathcal{I}_\nu$.

Consider a constraint $C \in \mathcal{C}_\nu^{\mathsf{out}}$. For partial assignments $b : \mathcal{T}_\nu \to [q]$ and $g : V \backslash S^\star \to [q]$, define $C(\rho_\nu \cup b \cup g) \in \{0, 1\}$ to be 1 iff C is satisfied by $\rho_\nu \cup b \cup g$. Since C only contains variables from $S_{\rho_\nu} \cup \mathcal{T}_\nu \cup (V \backslash S^\star)$, we have that $C(\rho_\nu \cup b \cup g)$ is well defined.

Define $\mathsf{score}^b : [q]^{V \backslash S^\star} \to \mathbb{R}$ by

$$\mathsf{score}^b(g) = \sum_{C \in \mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) \cdot C(\rho_\nu \cup b \cup g).$$

In words, $\mathsf{score}^{(b)}(g)$ is the total weight of constraints in $\mathcal{C}_\nu^{\mathsf{out}}$ satisfied by setting S_{ρ_ν} according to ρ_ν, setting $\mathcal{T}_\nu$ to b, and setting $V \backslash S^*$ according to g.

Note that for all $C \in \mathcal{C}_\nu^{\mathsf{out}}$, $\mathbf{E}_g[\sum_{b : \mathcal{T}_\nu \to [q]} C(\rho_\nu \cup b \cup g)] \geqslant \frac{1}{(qw)^{w - |\mathcal{T}_\nu|}}$. This follows since C is an active constraint given ρ_ν, and involves all variables from $\mathcal{T}_\nu$; hence there exists an assignment b to $\mathcal{T}_\nu$ and an assignment for at most $w - |\mathcal{T}_\nu|$ variables from constraint C in $V \backslash S^\star$ such that C is satisfied. Since, g is a smooth distribution, this particular assignment to $w - |\mathcal{T}_\nu|$ in $V \backslash S^\star$ is

sampled with probability at least $\frac{1}{(qw)^{w-|\mathcal{T}_\nu|}}$. Hence, for this particular choice of b, C is satisfied with probability at least $\frac{1}{(qw)^{w-|\mathcal{T}_\nu|}}$. Thus:

$$\sum_{b:\mathcal{T}_\nu\to[q]} \mathbf{E}_g\left[\mathsf{score}^b(g)\right] = \sum_{b:\mathcal{T}_\nu\to[q]} \mathbf{E}_g\left[\sum_{C\in\mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C)\cdot C(\rho_\nu\cup b\cup g)\right]$$
$$= \sum_{C\in\mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C)\cdot \mathbf{E}_g\left[\sum_{b:\mathcal{T}_\nu\to[q]} C(\rho_\nu\cup b\cup g)\right] \geqslant \frac{1}{(qw)^{w-|\mathcal{T}_\nu|}} \sum_{C\in\mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C).$$

Thus there exists $b:\mathcal{T}_\nu\to[q]$ such that

$$\mathbf{E}_g[\mathsf{score}^b(g)] \geqslant \frac{1}{q^{|\mathcal{T}_\nu|}}\cdot\frac{1}{(qw)^{w-|\mathcal{T}_\nu|}} \sum_{C\in\mathcal{C}_\nu^{\mathsf{out}}} \mathcal{W}_\ell(C) \geqslant \frac{1}{4}\cdot\frac{1}{(qw)^w}\cdot \mathsf{activedegree}_{\mathcal{T}_\nu},$$

where the last inequality follows by Lemma 4.7.

Fix this particular b for which the above inequality holds. We are going to use McDiarmid's inequality to show the concentration of $\mathsf{score}^b(g)$ around its mean. Since ν is typeC, from Lemma 4.7, we know that for every vertex $v\in V\backslash S^\star$, changing g on just v can change the value of $\mathsf{score}^b(g)$ by at most $c_v \stackrel{\text{def}}{=} \mathsf{activedegree}_{\rho_\nu}(\mathcal{T}_\nu\cup\{v\},\mathcal{W}_\ell) \leqslant \frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w}$. Thus by McDiarmid's inequality (Lemma D.1),

$$\begin{aligned}
\Pr_g[g \text{ is not } \mathsf{Cgood} \text{ for } \nu] &\leqslant \Pr_g\left[\mathsf{score}^b(g) < \frac{1}{8\cdot(qw)^w}\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}\right]\\
&\leqslant \Pr_g\left[|\mathsf{score}^b(g) - \mathbf{E}_g[\mathsf{score}^b(g)]| > \frac{1}{8\cdot(qw)^w}\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}\right]\\
&\leqslant 2\cdot\exp\left(\frac{-2\cdot\mathsf{activedegree}^2_{\mathcal{T}_\nu}}{64(qw)^{2w}\cdot\sum_{v\in V\backslash S^\star} c_v^2}\right).\\
&\leqslant 2\cdot\exp\left(\frac{-2\cdot\mathsf{activedegree}^2_{\mathcal{T}_\nu}}{64(qw)^{2w}\cdot(\max_v c_v)\cdot\sum_{v\in V\backslash S^\star} c_v}\right)\\
&\leqslant 2\cdot\exp\left(\frac{-2\cdot\mathsf{activedegree}^2_{\mathcal{T}_\nu}}{64(qw)^{2w}\cdot(\max_v c_v)\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}}\right)\\
&\leqslant 2\cdot\exp\left(\frac{-2\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}}{64(qw)^{2w}\cdot(\max_v c_v)}\right)\\
&\leqslant 2\cdot\exp\left(\frac{-2\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}}{64(qw)^{2w}\cdot(\frac{\mathsf{activedegree}_{\mathcal{T}_\nu}}{(4qwtk)^w})}\right)\\
&= 2\cdot\exp\left(\frac{-2\cdot(4qwtk)^w}{64\cdot(qw)^{2w}}\right) \leqslant 2\cdot\exp\left(\frac{-tk}{8qw}\right).
\end{aligned}$$

■

For a high variance instance ℓ, let $\nu_1^\ell,\ldots,\nu_t^\ell$ be the t nodes with $\mathcal{I}_\nu=\ell$ which lie on the path from the root to $\nu^\star$, numbered in order of their appearance on the path from the root to $\nu^\star$. Set

$\mathsf{finalwt}_\ell = \mathsf{activedegree}_{\rho^\star}(\mathcal{W}_\ell)$. This is the active degree left over in instance ℓ after the restriction $\rho^\star$.

Lemma 4.9 *For every high variance instance $\ell \in [k]$ and for each $i \leqslant [t/2]$,*

$$\mathsf{activedegree}_{\nu_i^\ell} \geqslant \gamma \cdot (1-\gamma)^{-t/2} \cdot \mathsf{finalwt}_\ell \geqslant 100 \cdot (qw)^w \cdot (4qwtk)^{w^2} \cdot \mathsf{finalwt}_\ell.$$

We skip the proof of this lemma. The first inequality is identical to the second part of Lemma 5.6, and the second inequality follows from the choice of t.

4.5.3 Putting everything together

We now show that when ν is taken to equal $\nu^\star$ in Step 5, then with high probability over the choice of $g_{\nu^\star}$ in Step 5(a) there is a setting of h in Step 5(b) such that $\min_{\ell \in [k]} \mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{q^{w-1}} - \varepsilon) \cdot c_\ell$.

Theorem 4.10 *Suppose the algorithm* Sim-MaxConjSAT *is given as inputs $\varepsilon > 0$, k simultaneous weighted* Max-w-ConjSAT$_q$ *instances $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on n variables, and target objective value $c_1, \ldots, c_k$ with the guarantee that there exists an assignment $f^\star$ such that for each $\ell \in [k]$, we have $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$. Then, the algorithm runs in $2^{O(k^4/\varepsilon^2 \log(k/\varepsilon))} \cdot \mathrm{poly}(n)$ time, and with probability at least 0.9, outputs an assignment f such that for each $\ell \in [k]$, we have, $\mathsf{val}(f, \mathcal{W}_\ell) \geqslant \left(\frac{1}{q^{w-1}} - \varepsilon\right) \cdot c_\ell$.*

Proof: Consider the case when ν is taken to equal $\nu^\star$ in Step 5. By Lemma 4.5, with probability at least $1 - k\delta_0$ over the random choices of $g_{\nu^\star}$, we have that for *every* low variance instance $\ell \in [k]$, $\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{q^{w-1}} - \frac{\varepsilon}{2}) \cdot c_\ell$. By Lemma 4.8 and a union bound, with probability at least $1 - \frac{t}{2} \cdot k \cdot 2e^{-tk/8qw} \geqslant 1 - \delta_0$ over the choice of $g_{\nu^\star}$, for every high variance instance ℓ and for every typeC node ν_i^ℓ, $i \in [t/2]$, we have that $g_{\nu^\star}$ is Cgood for ν_i^ℓ. Thus with probability at least $1 - (k+1)\delta_0$, both these events occur. Henceforth we assume that both these events occur in Step 5(a) of the algorithm.

Our next goal is to show that there exists a partial assignment $h : S^\star \to [q]$ such that

1. For every instance $\ell \in [k]$, $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell)$
2. For every high variance instance $\ell \in [k]$, $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot 10 \cdot \mathsf{finalwt}_\ell$.

Before giving a proof of the existence of such an h, we show that this completes the proof of the theorem. We claim that when the partial assignment h guaranteed above is considered in the Step 5(b) in the algorithm, we obtain an assignment with the required approximation guarantees.

For every low variance instance $\ell \in [k]$, since we started with $\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant (\frac{1}{q^{w-1}} - \frac{\varepsilon}{2}) \cdot c_\ell$, property 1 above implies that every low variance instance $\mathsf{val}(h \cup g_{\nu^\star}) \geqslant (\frac{1}{q^{w-1}} - \varepsilon) \cdot c_\ell$. For every high variance instance $\ell \in [k]$, since $h^\star = f^\star|_S$,

$$\mathsf{val}(h^\star \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \mathsf{val}(f^\star, \mathcal{W}_\ell) - \mathsf{activedegree}_{\rho^\star}(\mathcal{W}_\ell) \geqslant c_\ell - \mathsf{finalwt}_\ell.$$

Combining this with properties 1 and 2 above, we get,

$$\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(1 - \frac{\varepsilon}{2}\right) \cdot \max\{c_\ell - \mathsf{finalwt}_\ell, 10 \cdot \mathsf{finalwt}_\ell\} \geqslant \frac{10}{11}\left(1 - \frac{\varepsilon}{2}\right) \cdot c_\ell.$$

Thus, for all instances $\ell \in [k]$, we get $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant \left(\frac{1}{q^{w-1}} - \frac{\varepsilon}{2}\right) \cdot c_\ell$.

Now, it remains to show the existence of such an h by giving a procedure for constructing h by perturbing $h^\star$ (Note that this procedure is only part of the analysis). For nodes ν, ν' in the tree, let us write $\nu \prec \nu'$ if ν is an ancestor of ν', and we also say that ν' is "deeper" than ν.

Constructing h:

1. Initialize $H \subseteq [k]$ to be the set of high variance instances.
2. Let $N_0 = \{\nu_i^\ell \mid \ell \in H, i \in [t/2]\}$. Note that N is a chain in the tree (since all the elements of N are ancestors of $\nu^\star$). Since every $\nu \in N$ is an ancestor of $\nu^\star$, we have $h_{\rho_\nu} = h^\star|_{S_{\rho_\nu}}$.
3. Initialize $D = \varnothing$, $N = N_0$, $h = h^\star$.
4. During the procedure, we will be changing the assignment h, and removing elements from N. We will always maintain the following two invariants:
   - $|N| > \frac{t}{4}$.
   - For every $\nu \in N$, $h|_{S_{\rho_\nu}} = h^\star|_{S_{\rho_\nu}}$.
5. While $|D| \neq |H|$ do:
   (a) Let
   $$B = \left\{ v \in V \mid \exists \ell \in [k] \text{ with } \sum_{C \in \mathcal{C}, C \ni v} \mathcal{W}_\ell(C) \cdot C(h \cup g_{\nu^\star}) \geqslant \frac{\varepsilon}{2wk} \mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \right\}.$$
   Note that $|B| \leqslant \frac{2w^2k^2}{\varepsilon} < \frac{t}{4}$.
   (b) Let $\nu \in N$ be the deepest element of N for which: $\mathcal{T}_\nu \cap B = \varnothing$.
   Such a ν exists because:
   - $|N| > \frac{t}{4} > |B|$, and
   - there are at most $|B|$ nodes ν for which $\mathcal{T}_\nu \cap B \neq \varnothing$ (since $\mathcal{T}_\nu$ are all disjoint for distinct ν).

   (c) Let $\ell \in H$ and $i \in [t/2]$ be such that $\nu = \nu_i^\ell$. Let $\rho = \rho_\nu$. We will now modify the assignment h for variables in $\mathcal{T}_\nu$ to guarantee that $\mathsf{val}(h \cup g_{\nu^\star}, \mathcal{W}_\ell) \geqslant 10 \cdot \mathsf{finalwt}_\ell$. The procedure depends on whether ν is typeAB or typeC.
      i. If ν is typeAB, then we know that $\mathsf{backward}_\nu \geqslant \frac{1}{2} \cdot \mathsf{activedegree}_{\mathcal{T}_\nu}$.
      The second invariant tells us that $\rho = h^\star|_{S_\rho} = h|_{S_\rho}$. Thus we have:
      $$\begin{aligned} \mathsf{backward}_\nu &= \sum_{C \in \mathcal{C}_\nu^{\mathsf{backward}}} \mathcal{W}_\ell(C) \\ &= \sum_{C \subseteq S_\rho \cup \mathcal{T}_\nu, C \supseteq \mathcal{T}_\nu, C \in \mathsf{Active}(\rho)} \mathcal{W}_\ell(C) \\ &= \sum_{C \subseteq S_\rho \cup \mathcal{T}_\nu, C \supseteq \mathcal{T}_\nu, C \in \mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C). \end{aligned}$$

This implies that we can modify the assignment h on the variables $\mathcal{T}_\nu$ such that after the modification, the weights of satisfied backward constraints is:

$$\begin{aligned}\sum_{C\subseteq S_\rho\cup\mathcal{T}_\nu, C\supseteq\mathcal{T}_\nu, C\in\mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C)C(h) &\geqslant \frac{1}{q^w}\sum_{C\subseteq S_\rho\cup\mathcal{T}_\nu, C\supseteq\mathcal{T}_\nu, C\in\mathsf{Active}(h|_{S_\rho})} \mathcal{W}_\ell(C)\\ &= \frac{1}{q^w}\cdot \mathsf{backward}_\nu\\ &\geqslant \frac{1}{2q^w}\cdot \mathsf{activedegree}_{\mathcal{T}_\nu}\\ &\geqslant 10\cdot\mathsf{finalwt}_\ell.\end{aligned}$$

where the $\frac{1}{q^w}$ factor in the first inequality appears because there could be as many as q^w possible assignments to variables in $\mathcal{T}_\nu$, and the last inequality holds because of Observation 4.6 and Lemma 4.9. In particular, after making this change, we have $\mathsf{val}(h\cup g_{\nu^\star}, \mathcal{W}_\ell)\geqslant 10\cdot\mathsf{finalwt}_\ell$.

ii. If ν is typeC, then we know that g is Cgood for ν. Thus, by the definition of Cgood, we can choose a setting of $\mathcal{T}_\nu$ so that at least a total of $\frac{1}{8\cdot(qw)^w}\cdot\mathsf{activedegree}_{\mathcal{T}_\nu}\geqslant 10\cdot\mathsf{finalwt}_\ell$ $\mathcal{W}_\ell$-weight constraints between $\mathcal{T}_\nu$ and $V\backslash S^\star$ is satisfied. After this change, we have $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_\ell)\geqslant 10\cdot\mathsf{finalwt}_\ell$.

In both the above cases, we only changed the value of h at the variables $\mathcal{T}_\nu$. Since $\mathcal{T}_\nu\cap B=\varnothing$, we have that for every $j\in[k]$, the new value $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_j)$ is at least $\left(1-\frac{\varepsilon}{2k}\right)$ times the old value $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_j)$.

(d) Set $D=D\cup\{\ell\}$.

(e) Set $N=\{\nu_i^\ell \mid \ell\in H\backslash D, i\leqslant[t/2], \nu_i^\ell\prec\nu\}$.
Observe that $|N|$ decreases in size by at most $\frac{t}{2}+|B|$. Thus, if $D\neq H$, we have

$$\begin{aligned}|N| &\geqslant |N_0|-|D|\cdot\frac{t}{2}-|D||B|\\ &= |H|\cdot\frac{t}{2}-|D|\cdot\frac{t}{2}-|D||B|\\ &\geqslant \frac{t}{2}-k|B| > \frac{t}{4}\end{aligned}$$

Also observe that we only changed the values of h at the variables $\mathcal{T}_\nu$. Thus for all $\nu'\preceq\nu$ (i.e $\nu'\in N$), we still have the property that $h|_{S_{\rho_{\nu'}}}=h^\star|_{S_{\rho_{\nu'}}}$.

For each high variance instance $\ell\in[k]$, in the iteration where ℓ gets added to the set D, the procedure ensures that at the end of the iteration $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_\ell)\geqslant 10\cdot\mathsf{finalwt}_\ell$.

Moreover, at each step we reduced the value of each $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_\ell)$ by at most $\frac{\varepsilon}{2k}$ fraction of its previous value. Thus, at the end of the procedure, for every $\ell\in[k]$, the value has decreased at most by a multiplicative factor of $\left(1-\frac{\varepsilon}{2k}\right)^k\geqslant\left(1-\frac{\varepsilon}{2}\right)$. Thus, for every $\ell\in[k]$, we get $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_\ell)\geqslant\left(1-\frac{\varepsilon}{2}\right)\cdot\mathsf{val}(h^\star\cup g_{\nu^\star},\mathcal{W}_\ell)$, and for every high variance instance $\ell\in[k]$, we have $\mathsf{val}(h\cup g_{\nu^\star},\mathcal{W}_\ell)\geqslant\left(1-\frac{\varepsilon}{2}\right)\cdot 10\cdot\mathsf{finalwt}_\ell$. This proves the two properties of h that we set out to prove.

Running time : Running time of the algorithm is $2^{O(kt)}\cdot\mathrm{poly}(n)$ which is $2^{O(k^4/\varepsilon^2\log(k/\varepsilon^2))}\cdot\mathrm{poly}(n)$. ■

5 Simultaneous Max-w-SAT

In this section, we give our algorithm for simultaneous Max-w-SAT. The algorithm follows the basic paradigm from Max-2-AND and Max-CSP, but does not require a tree of evolutions (only a set of influential variables), and uses an LP to boost the Pareto approximation factor to $\left(\frac{3}{4} - \varepsilon\right)$.

5.1 Preliminaries

Let V be a set of n Boolean variables. Define $\mathcal{C}$ to be the set of all possible w-SAT constraints on the n variable set V. A Max-w-SAT instance is then described by a weight function $\mathcal{W} : \mathcal{C} \to \mathbb{R}_{\geqslant 0}$ (here $\mathcal{W}(C)$ denotes the weight of the constraint C). We will assume that $\sum_{C \in \mathcal{C}} \mathcal{W}(C) = 1$.

We say $v \in C$ if the variable v appears in the constraint C. For a constraint C, let C^+ (resp. C^-) denote the set of variables $v \in V$ that appear unnegated (resp. negated) in the constraint C.

Let $f : V \to \{0,1\}$ be an assignment. For a constraint $C \in \mathcal{C}$, define $C(f)$ to be 1 if the constraint C is satisfied by the assignment f, and define $C(f) = 0$ otherwise. Then, we have the following expression for $\mathsf{val}(f, \mathcal{W})$:

$$\mathsf{val}(f, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{C \in \mathcal{C}} \mathcal{W}(C) \cdot C(f).$$

5.1.1 Active Constraints

Our algorithm will maintain a small set $S \subseteq V$ of variables, for which we will try all assignments by brute-force, and then use a randomized rounding procedure for a linear program to obtain an assignment for $V \backslash S$. We now introduce some notation for dealing with this.

Let $S \subseteq V$. We say a constraint $C \in \mathcal{C}$ is active given S if at least one of the variables of C is in $V \backslash S$. We denote by $\mathsf{Active}(S)$ the set of constraints from $\mathcal{C}$ which are active given S. For two constraints $C_1, C_2 \in \mathcal{C}$, we say $C_1 \sim_S C_2$ if they share a variable that is contained in $V \backslash S$. Note that if $C_1 \sim_S C_2$, then C_1, C_2 are both in $\mathsf{Active}(S)$. For two partial assignments $f_1 : S \to \{0,1\}$ and $f_2 : V \backslash S \to \{0,1\}$, let $f = f_1 \cup f_2$ is an assignment $f : V \to \{0,1\}$ such that $f(x) = f_1(x)$ if $x \in S$ otherwise $f(x) = f_2(x)$.

Define the active degree of a variable $v \in V \backslash S$ given S by:

$$\mathsf{activedegree}_S(v, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{C \in \mathsf{Active}(S), C \ni v} \mathcal{W}(C).$$

We then define the active degree of the whole instance $\mathcal{W}$ given S:

$$\mathsf{activedegree}_S(\mathcal{W}) \stackrel{\text{def}}{=} \sum_{v \in V \backslash S} \mathsf{activedegree}_S(v, \mathcal{W}).$$

For a partial assignment $h : S \to \{0,1\}$, we define

$$\mathsf{val}(h, \mathcal{W}) \stackrel{\text{def}}{=} \sum_{\substack{C \in \mathcal{C} \\ C \notin \mathsf{Active}(S)}} \mathcal{W}(C) \cdot C(h).$$

Thus, for an assignment $g : V \backslash S \to \{0,1\}$, to the remaining set of variables, we have the equality:

$$\mathsf{val}(h \cup g, \mathcal{W}) - \mathsf{val}(h, \mathcal{W}) = \sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C) \cdot C(h \cup g).$$

5.1.2 LP Rounding

Let $h : S \to \{0,1\}$ be a partial assignment. We will use the Linear Program $\mathsf{MAXwSAT\text{-}LP}_1(h)$ to complete the assignment to $V\backslash S$. For MAX-2-SAT, Goemans and Williamson [GW93] showed, via a rounding procedure, that this LP can be used to give a 3/4 approximation. However, as in MAX-2-AND, we will be using the rounding procedure due to Trevisan [Tre98] that also gives a 3/4 approximation for MAX-w-SAT, because of its smoothness properties.

Let $\vec{t}, \vec{z}$ be a feasible solution to the LP $\mathsf{MAXwSAT\text{-}LP}_1(h)$. Let $\mathsf{smooth}(\vec{t})$ denote the map $p : V\backslash S \to [0,1]$ given by: $p(v) = \frac{1}{4} + \frac{t_v}{2}$. Note that $p(v) \in [1/4, 3/4]$ for all v.

Theorem 5.1 *Let $h : S \to \{0,1\}$ be a partial assignment.*

1. *For every $g_0 : V\backslash S \to \{0,1\}$, there exist $\vec{t}$, $\vec{z}$ satisfying $\mathsf{MAXwSAT\text{-}LP}_1(h)$ such that for every* MAX-w-SAT *instance $\mathcal{W}$:*
$$\sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C = \mathsf{val}(g_0 \cup h, \mathcal{W}).$$

2. *Suppose $\vec{t}, \vec{z}$ satisfy $\mathsf{MAXwSAT\text{-}LP}_1(h)$. Let $p = \mathsf{smooth}(\vec{t})$. Then for every* MAX-$w$-SAT *instance $\mathcal{W}$:*
$$\mathop{\mathbf{E}}_{g}[\mathsf{val}(h \cup g, \mathcal{W})] \geqslant \frac{3}{4} \cdot \sum_{C \in \mathcal{C}} \mathcal{W}(C) z_C,$$
where $g : V\backslash S \to \{0,1\}$ is such that each $g(v)$ is sampled independently with $\mathbf{E}[g(v)] = p(v)$.

Proof: The first part is identical to the first part of Lemma 3.2. For the second part. Let $\mathcal{W}$ be any instance of MAX-w-SAT. Let $g : V\backslash S \to \{0,1\}$ be sampled as follows: independently for each $v \in V\backslash S$, $g(v)$ is sampled from $\{0,1\}$ such that $\mathbf{E}[g(v)] = p(v)$. We need to show that

$$\begin{aligned}
\mathop{\mathbf{E}}_{g}[\mathsf{val}(h \cup g, \mathcal{W})] &= \sum_{C \in \mathcal{C}\backslash \mathsf{Active}(S)} \mathcal{W}(C) C(h) + \mathbf{E}\left[\sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C) C(\rho \cup g)\right] \\
&\geqslant \frac{3}{4} \cdot \sum_{C \in \mathcal{C}\backslash \mathsf{Active}(S)} \mathcal{W}(C) z_C + \frac{3}{4} \cdot \sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C) z_C
\end{aligned}$$

For $C \in \mathcal{C}\backslash\mathsf{Active}(S)$, it is easy to verify that if $z_C > 0$, we must have $C(h) = 1$. For $C \in \mathsf{Active}(S)$ the following claim gives us the required inequality:

Claim 5.2 *For $C \in \mathsf{Active}(S)$, $\mathbf{E}[C(h \cup g)] \geqslant \frac{3}{4} \cdot z_C$.*

Proof: The claim is true if C is satisfied by h. Consider a clause C which contains l active variables but not satisfied by partial assignment h. Under the smooth rounding, we have

$$\begin{aligned}
\mathbf{E}[C(h \cup g)] = \Pr[C \text{ is satisfied by } h \cup g] &= 1 - \left(\prod_{v \in C^+, v \in V\backslash S} \frac{3}{4} - \frac{t_v}{2}\right) \cdot \left(\prod_{v \in C^-, v \in V\backslash S} \frac{3}{4} - \frac{1 - t_v}{2}\right) \\
&\geqslant 1 - \left(\frac{3}{4} - \frac{\sum_{v \in C^+, v \in V\backslash S} t_v + \sum_{v \in C^-, v \in V\backslash S} (1 - t_v)}{2l}\right)^l \\
&\geqslant 1 - \left(\frac{3}{4} - \frac{z_C}{2l}\right)^l \geqslant \frac{3}{4} \cdot z_C,
\end{aligned}$$

where first inequality follows from AM-GM inequality. For any integer $l \geqslant 1$, the last inequality follows by noting that for a function $f(x) = 1 - \left(\frac{3}{4} - \frac{x}{2l}\right)^l - \frac{3}{4} \cdot x$, $f(0) \geqslant 0, f(1) \geqslant 0$ along with the fact the the function has no local minima in $(0,1)$. ■

■

5.2 Random Assignments

We now give a sufficient condition for the value of a MAX-w-SAT instance to be highly concentrated under a *sufficiently smooth* independent random assignment to the variables of $V \backslash S$ (This smooth distribution will come from the rounding algorithm for the LP). When the condition does not hold, we will get a variable of high active degree.

Let $S \subseteq V$, and let $h : S \to \{0,1\}$ be an arbitrary partial assignment to S. Let $p : V \backslash S \to [0,1]$ be such that $p(v) \in [1/4, 3/4]$ for each $v \in V \backslash S$. Consider the random assignment $g : V \backslash S \to \{0,1\}$, where for each $v \in V \backslash S$, $g(v) \in \{0,1\}$ is sampled independently with $\mathbf{E}[g(v)] = p(v)$. Define the random variable

$$Y \stackrel{\text{def}}{=} \mathsf{val}(h \cup g, \mathcal{W}) - \mathsf{val}(h, \mathcal{W}) = \sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C) \cdot C(h \cup g).$$

The random variable Y measures the contribution of active constraints to the instance $\mathcal{W}$.

We now define two quantities depending only on S (and importantly, not on h), which will be useful in controlling the expectation and variance of Y. The first quantity is an upper bound on $\mathbf{Var}[Y]$:

$$\mathsf{Uvar} \stackrel{\text{def}}{=} \sum_{C_1 \sim_S C_2} \mathcal{W}(C_1)\mathcal{W}(C_2).$$

The second quantity is a lower bound on $\mathbf{E}[Y]$:

$$\mathsf{Lmean} \stackrel{\text{def}}{=} \frac{1}{4} \cdot \sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C).$$

Lemma 5.3 *Let $S \subseteq V$ be a subset of variables and $h : S \to \{0,1\}$ be an arbitrary partial assignment to S. Let $p, Y, \mathsf{Uvar}, \mathsf{Lmean}$ be as above.*

1. *If $\mathsf{Uvar} \leqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{Lmean}^2$, then $\Pr[Y < (1 - \varepsilon_0)\mathbf{E}[Y]] < \delta_0$.*
2. *If $\mathsf{Uvar} \geqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{Lmean}^2$, then there exists $v \in V \backslash S$ such that*

$$\mathsf{activedegree}_S(v, \mathcal{W}) \geqslant \frac{1}{16w^2} \varepsilon_0^2 \delta_0 \cdot \mathsf{activedegree}_S(\mathcal{W}).$$

The crux of the proof is that independent of the assignment $h : S \to \{0,1\}$, $\mathbf{E}[Y] \geqslant \mathsf{Lmean}$ and $\mathbf{Var}(Y) \leqslant \mathsf{Uvar}$ (this crucially requires that the rounding is independent and *smooth*, *i.e.*, $p(v) \in [1/4, 3/4]$ for all v; this is why we end up using Trevisan's rounding procedure in Theorem 5.1). The first part is then a simple application of the Chebyshev inequality. For the second part, we use the assumption that Uvar is large, to deduce that there exists a constraint C such that the total weight of constraints that share a variable from $V \backslash S$ with C, *i.e.*, $\sum_{C_2 \sim_S C} \mathcal{W}(C_2)$, is large. It then follows that at least one variable $v \in C$ must have large activedegree given S.

Proof: We first prove that $\mathbf{Var}(Y) \leqslant \mathsf{Uvar}$. Recall that the indicator variable $C(h \cup g)$ denotes whether a constraint C is satisfied by the assignment $h \cup g$, and note that:

$$Y = \sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C) \cdot C(h \cup g).$$

Thus, the variance of Y is given by

$$\begin{aligned}\mathbf{Var}(Y) &= \sum_{C_1, C_2 \in \mathsf{Active}(S)} \mathcal{W}(C_1)\mathcal{W}(C_2) \cdot (\mathbf{E}[C_1(h \cup g)C_2(h \cup g)] - \mathbf{E}[C_1(h \cup g)]\mathbf{E}[C_2(h \cup g)]) \\ &\leqslant \sum_{C_1 \sim_S C_2} \mathcal{W}(C_1)\mathcal{W}(C_2) = \mathsf{Uvar},\end{aligned}$$

where the inequality holds because $\mathbf{E}[C_1(h \cup g)C_2(h \cup g)] - \mathbf{E}[C_1(h \cup g)]\mathbf{E}[C_2(h \cup g)] \leqslant 1$ for all C_1, C_2, and $\mathbf{E}[C_1(h \cup g)C_2(h \cup g)] - \mathbf{E}[C_1(h \cup g)]\mathbf{E}[C_2(h \cup g)] = 0$ unless $C_1 \sim_S C_2$ because the rounding is performed independently for all the variables.

Moreover, since $p(v) \in [1/4, 3/4]$ for all v, we get that $\mathbf{E}[C(h \cup g)] \geqslant 1/4$ for all $C \in \mathsf{Active}(S)$. Thus, we have $\mathbf{E}[Y] \geqslant \mathsf{Lmean}$. Given this, the first part of the lemma easily follows from Chebyshev's inequality:

$$\Pr[Y < (1 - \varepsilon_0)\mathbf{E}[Y]] \leqslant \frac{\mathbf{Var}(Y)}{\varepsilon_0^2(\mathbf{E}[Y])^2} \leqslant \frac{\mathsf{Uvar}}{\varepsilon_0^2 \mathsf{Lmean}^2} \leqslant \delta_0.$$

For the second part of the lemma, we have:

$$\begin{aligned}\delta_0\varepsilon_0^2\mathsf{Lmean}^2 < \mathsf{Uvar} &= \sum_{C_1 \sim_S C_2} \mathcal{W}(C_1)\mathcal{W}(C_2) \\ &\leqslant \sum_{C_1 \in \mathsf{Active}(S)} \mathcal{W}(C_1) \sum_{C_2 \sim_S C_1} \mathcal{W}(C_2) \\ &\leqslant \left(\sum_{C_1 \in \mathsf{Active}(S)} \mathcal{W}(C_1) \right) \cdot \max_{C \in \mathsf{Active}(S)} \sum_{C_2 \sim_S C} \mathcal{W}(C_2) \\ &= 4 \cdot \mathsf{Lmean} \cdot \max_{C \in \mathsf{Active}(S)} \sum_{C_2 \sim_S C} \mathcal{W}(C_2)\ .\end{aligned}$$

Thus, there exists a constraint $C \in \mathsf{Active}(S)$ such that:

$$\sum_{C_2 \sim_S C} \mathcal{W}(C_2) \geqslant \frac{1}{4} \cdot \delta_0\varepsilon_0^2 \cdot \mathsf{Lmean} \geqslant \frac{1}{16w}\delta_0\varepsilon_0^2 \cdot \mathsf{activedegree}_S(\mathcal{W}), \tag{2}$$

where we used the fact that $\mathsf{Lmean} = \frac{1}{4} \cdot (\sum_{C \in \mathsf{Active}(S)} \mathcal{W}(C)) \geqslant \frac{1}{4w} \cdot \mathsf{activedegree}_S(\mathcal{W})$, since we are counting the weight of a constraint at most w times in the expression $\mathsf{activedegree}_S(\mathcal{W})$. Finally, the LHS of equation (2) is at most $\sum_{u \in C \cap (V \setminus S)} \mathsf{activedegree}_S(u, \mathcal{W})$. Thus, there is some $u \in V \setminus S$ with:

$$\mathsf{activedegree}_S(u, \mathcal{W}) \geqslant \frac{1}{16w^2}\delta_0\varepsilon_0^2 \cdot \mathsf{activedegree}_S(\mathcal{W}).$$

■

Input: k instances of MAX-w-SAT $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on the variable set V, target objective values $c_1, \ldots, c_k$, and $\varepsilon > 0$.
Output: An assignment to V.
Parameters: $\delta_0 = \frac{1}{10k}$, $\varepsilon_0 = \frac{\varepsilon}{2}$, $\gamma = \frac{\varepsilon_0^2 \delta_0}{16 w^2}$, $t = \frac{2k}{\gamma} \cdot \log\left(\frac{11}{\gamma}\right)$.

1. Initialize $S \leftarrow \varnothing$.
2. For each instance $\ell \in [k]$, initialize $\mathsf{count}_\ell \leftarrow 0$ and $\mathsf{flag}_\ell \leftarrow$ TRUE.
3. Repeat the following until for every $\ell \in [k]$, either $\mathsf{flag}_\ell =$ FALSE or $\mathsf{count}_\ell = t$:
   (a) For each $\ell \in [k]$, compute $\mathsf{Uvar}_\ell = \sum_{C_1 \sim_S C_2} \mathcal{W}_\ell(C_1)\mathcal{W}_\ell(C_2)$.
   (b) For each $\ell \in [k]$, compute $\mathsf{Lmean}_\ell = \frac{1}{4}\sum_{C \in \mathsf{Active}(S)} \mathcal{W}_\ell(C)$.
   (c) For each $\ell \in [k]$, if $\mathsf{Uvar}_\ell \geqslant \delta_0 \varepsilon_0^2 \cdot \mathsf{Lmean}_\ell^2$, then set $\mathsf{flag}_\ell =$ TRUE, else set $\mathsf{flag}_\ell =$ FALSE.
   (d) Choose any $\ell \in [k]$, such that $\mathsf{count}_\ell < t$ AND $\mathsf{flag}_\ell =$ TRUE (if any):
      i. Find a variable $v \in V$ such that $\mathsf{activedegree}_S(v, \mathcal{W}_\ell) \geqslant \gamma \cdot \mathsf{activedegree}_S(\mathcal{W}_\ell)$.
      ii. Set $S \leftarrow S \cup \{v\}$. We say that v was brought into S because of instance ℓ.
      iii. Set $\mathsf{count}_\ell \leftarrow \mathsf{count}_\ell + 1$.
4. For each partial assignment $h_0 : S \to \{0,1\}$:
   (a) If there is a feasible solution $\vec{t}, \vec{z}$ to the LP in Figure 10, set $p = \mathsf{smooth}(\vec{t})$. If not, return to Step 4. and proceed to the next h_0.
   (b) Define $g : V \backslash S \to \{0,1\}$ by independently sampling $g(v) \in \{0,1\}$ with $\mathbf{E}[g(v)] = p(v)$, for each $v \in V \backslash S$.
   (c) For each $h : S \to \{0,1\}$, compute $\mathsf{out}_{h,g} = \min_{\ell \in [k]} \frac{\mathsf{val}(h \cup g, \mathcal{W}_\ell)}{c_\ell}$. If $c_\ell = 0$ for some $\ell \in [k]$, we interpret $\frac{\mathsf{val}(h \cup g, \mathcal{W}_\ell)}{c_\ell}$ as $+\infty$.
5. Output the largest $\mathsf{out}_{h,g}$ seen, and the assignment $h \cup g$.

Figure 8: Algorithm SIM-MAXWSAT for approximating weighted simultaneous MAX-w-SAT

$$
\begin{array}{rcll}
\sum_{v \in C^+} t_v + \sum_{v \in C^-} (1 - t_v) & \geqslant & z_C & \forall C \in \mathcal{C} \\
1 \geqslant z_C & \geqslant & 0 & \forall C \in \mathcal{C} \\
1 \geqslant t_v & \geqslant & 0 & \forall v \in V \backslash S \\
t_v & = & h_0(v) & \forall v \in S
\end{array}
$$

Figure 9: Linear program $\mathsf{MAXwSAT\text{-}LP}_1(h_0)$, for a given partial assignment $h_0 : S \to \{0,1\}$

$$
\begin{array}{r}
\sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) \cdot z_C \geqslant c_\ell \quad \forall \ell \in [k] \\
\vec{t}, \vec{z} \text{ satisfy } \mathsf{MAXwSAT\text{-}LP}_1(h_0).
\end{array}
$$

Figure 10: Linear program $\mathsf{MAXwSAT\text{-}LP}_2(h_0)$ for a given partial assignment $h_0 : S \to \{0,1\}$

5.3 Algorithm for Simultaneous MAX-w-SAT

In Figure 8, we give our algorithm for simultaneous MAX-w-SAT. The input to the algorithm consists of an integer $k \geqslant 1$, $\varepsilon > 0$, and k instances of MAX-w-SAT, specified by weight functions $\mathcal{W}_1, \ldots, \mathcal{W}_k$, and target objective values $c_1, \ldots, c_\ell$.

5.4 Analysis of Algorithm SIM-MAXWSAT

It is easy to see that the algorithm always terminates in polynomial time. Part 2 of Lemma 5.3 implies that that Step 3.(d)i always succeeds in finding a variable v. Next, we note that Step 3. always terminates. Indeed, whenever we find an instance $\ell \in [k]$ in Step 3.d such that $\mathsf{count}_\ell < t$ and $\mathsf{flag}_\ell = \textsc{True}$, we increment count_ℓ. This can happen only tk times before the condition $\mathsf{count}_\ell < t$ fails for all $\ell \in [k]$. Thus the loop must terminate within tk iterations.

Let $S^\star$ denote the final set S that we get at the end of Step 3. of SIM-MAXWSAT. To analyze the approximation guarantee of the algorithm, we classify instances according to how many vertices were brought into $S^\star$ because of them.

Definition 5.4 (Low and high variance instances) *At the completion of Step 3.d in Algorithm* SIM-MAXWSAT*, if $\ell \in [k]$ satisfies $\mathsf{count}_\ell = t$, we call instance ℓ a* high variance *instance. Otherwise we call instance ℓ a* low variance *instance.*

At a high level, the analysis will go as follows: First we analyze what happens when we give the optimal assignment to $S^\star$ in Step 4. For low variance instances, the fraction of the constraints staisfied by the LP rounding will concentrate around its expectation, and will give the desired approximation. For every high variance instance, we will see that many of its "heavy-weight" vertices were brought into $S^\star$, and we will use this to argue that we can satisfy a large fraction of the constraints from these high variance instances by suitably perturbing the optimal assignment to $S^\star$ to these "heavy-weight" vertices. It is crucial that this perturbation is carried out without significantly affecting the value of the low variance instances.

Let $f^\star : V \to \{0,1\}$ be an assignment such that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for each ℓ. Let $h^\star = f^\star|_{S^\star}$. Claim 1 from Theorem 5.1 implies that $\mathsf{MAXwSAT\text{-}LP}_2(h^\star)$ has a feasible solution. For low variance instances, by combining Theorem 5.1 and Lemma 5.3, we show that $\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$ is at least $(3/4 - \varepsilon/2) \cdot c_\ell$ with high probability.

Lemma 5.5 *Let $\ell \in [k]$ be any low variance instance. Let $\vec{t}, \vec{z}$ be a feasible solution to $\mathsf{MAXwSAT\text{-}LP}_2(h^\star)$. Let $p = \mathsf{smooth}(\vec{t})$. Let $g : V \backslash S^\star \to \{0,1\}$ be such that each $g(v)$ is sampled independently with $\mathbf{E}[g(v)] = p(v)$. Then the assignment $h^\star \cup g$ satisfies:*

$$\Pr_g \left[\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) \geqslant (3/4 - \varepsilon/2) \cdot c_\ell\right] \geqslant 1 - \delta_0.$$

Proof: Since ℓ is a *low variance* instance, $\mathsf{flag}_\ell = \textsc{False}$ when the algorithm terminates. Thus $\mathsf{Uvar}_\ell < \delta_0 \varepsilon_0^2 \cdot \mathsf{Lmean}_\ell^2$. Let $g : V \to \{0,1\}$ be the random assignment picked in Step 4.b. Define the random variable

$$Y_\ell \stackrel{\text{def}}{=} \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) - \mathsf{val}(h^\star, \mathcal{W}_\ell).$$

By Lemma 5.3, we know that with probability at least $1 - \delta_0$, we have $Y_\ell \geqslant (1 - \varepsilon_0)\mathbf{E}[Y_\ell]$. Thus, with probability at least $1 - \delta_0$, we have,

$$\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) = \mathsf{val}(h^\star, \mathcal{W}_\ell) + Y_\ell \geqslant \mathsf{val}(h^\star, \mathcal{W}_\ell) + (1 - \varepsilon_0)\mathbf{E}[Y_\ell]$$

$$\geqslant (1-\varepsilon_0)\cdot \mathbf{E}[\mathsf{val}(h^\star, \mathcal{W}_\ell) + Y_\ell] = (1-\varepsilon_0)\cdot \mathbf{E}[\mathsf{val}(h^\star \cup g, W_\ell)]$$
$$\geqslant 3/4 \cdot (1-\varepsilon_0)\cdot \sum_{C\in\mathcal{C}} \mathcal{W}_\ell(C) z_C \geqslant (3/4 - \varepsilon/2)\cdot c_\ell,$$

where the last two inequalities follow from Claim 2 in Theorem 5.1 and the constraints in MAXwSAT-LP$_2$ respectively. ■

Now we analyze the high variance instances. We prove the following lemma that proves that at the end of the algorithm, the activedegree of high variance instances is small, and is dominated by the activedegree of any variable that was included in S "early on".

Lemma 5.6 *For all high variance instances $\ell \in [k]$, we have*

1. $\mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell) \leqslant w(1-\gamma)^t$.
2. *For each of the first $t/2$ variables that were brought inside $S^\star$ because of instance ℓ, the total weight of constraints incident on each of that variable and totally contained inside $S^\star$ is at least $10 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$.*

The crucial observation is that when a variable u is brought into S because of an instance ℓ, the activedegree of u is at least a γ fraction of the total activedegree of instance ℓ. Thus, the activedegree of instance ℓ goes down by a multiplicative factor of $(1-\gamma)$. This immediately implies the first part of the lemma. For the second part, we use the fact that t is large, and hence the activedegree of early vertices must be much larger than the final activedegree of instance ℓ.

Proof: Consider any *high variance* instance $\ell \in [k]$. Initially, when $S = \varnothing$, we have $\mathsf{activedegree}_{\varnothing}(\mathcal{W}_\ell) \leqslant w$ since the weight of every constraint is counted at most w times, once for each of the 2 active variables of the constraint, and $\sum_{C\in\mathcal{C}} \mathcal{W}_\ell(C) = 1$. For every v, note that $\mathsf{activedegree}_{S_2}(v, \mathcal{W}_\ell) \leqslant \mathsf{activedegree}_{S_1}(v, \mathcal{W}_\ell)$ whenever $S_1 \subseteq S_2$.

Let u be one of the variables that ends up in $S^\star$ because of instance ℓ. Let S_u denote the set $S \subseteq S^\star$ just before u was brought into $S^\star$. When u is added to S_u, we know that $\mathsf{activedegree}_{S_u}(u, \mathcal{W}_\ell) \geqslant \gamma \cdot \mathsf{activedegree}_{S_u}(\mathcal{W}_\ell)$. Hence, $\mathsf{activedegree}_{S_u\cup\{u\}}(\mathcal{W}_\ell) \leqslant \mathsf{activedegree}_{S_u}(\mathcal{W}_\ell) - \mathsf{activedegree}_{S_u}(u, \mathcal{W}_\ell) \leqslant (1-\gamma)\cdot \mathsf{activedegree}_{S_u}(\mathcal{W}_\ell)$. Since t variables were brought into $S^\star$ because of instance ℓ, and initially $\mathsf{activedegree}_{\varnothing}(\mathcal{W}_\ell) \leqslant w$, we get $\mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell) \leqslant w(1-\gamma)^t$.

Now, let u be one of the first $t/2$ variables that ends up in $S^\star$ because of instance ℓ. Since at least $t/2$ variables are brought into $S^\star$ because of instance ℓ, after u, as above, we get $\mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell) \leqslant (1-\gamma)^{t/2}\cdot \mathsf{activedegree}_{S_u}(\mathcal{W}_\ell)$. Combining with $\mathsf{activedegree}_{S_u}(u, \mathcal{W}_\ell) \geqslant \gamma \cdot \mathsf{activedegree}_{S_u}(\mathcal{W}_\ell)$, we get $\mathsf{activedegree}_{S_u}(u, \mathcal{W}_\ell) \geqslant \gamma(1-\gamma)^{-t/2}\mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$, which is at least $11 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$, by the choice of parameters. Since any constraint incident on a vertex in $V\backslash S^\star$ contributes its weight to $\mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$, the total weight of constraints incident on u and totally contained inside $S^\star$ is at least $10 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$ as required. ■

We now describe a procedure PERTURB (see Figure 11) which takes $h^\star : S^\star \to \{0,1\}$ and $g : V\backslash S^\star \to \{0,1\}$, and produces a new $h : S^\star \to \{0,1\}$ such that for all (low variance as well as high variance) instances $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell)$ is not much smaller than $\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$, and furthermore, for all high variance instances $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell)$ is large. The procedure works by picking a special variable in $S^\star$ for every high variance instance and perturbing the assignment of $h^\star$ to these special variables. The crucial feature used in the perturbation procedure, which holds

Input: $h^\star : S^\star \to \{0,1\}$ and $g : V \backslash S^\star \to \{0,1\}$
Output: A perturbed assignment $h : S^\star \to \{0,1\}$.

1. Initialize $h \leftarrow h^\star$.
2. For $\ell = 1, \ldots, k$, if instance ℓ is a high variance instance case (i.e., $\mathsf{count}_\ell = t$), we pick a special variable $v_\ell \in S^\star$ associated to this instance as follows:
   (a) Let $B = \{v \in V \mid \exists \ell \in [k]$ with $\sum_{C \in \mathcal{C}, C \ni v} \mathcal{W}_\ell(C) \cdot C(h \cup g) \geqslant \frac{\varepsilon}{2k} \cdot \mathsf{val}(h \cup g, \mathcal{W}_\ell)\}$. Since the weight of each constraint is counted at most w times, we know that $|B| \leqslant \frac{2wk^2}{\varepsilon}$.
   (b) Let U be the set consisting of the first $t/2$ variables brought into $S^\star$ because of instance ℓ.
   (c) Since $t/2 > |B| + k$, there exists some $u \in U$ such that $u \notin B \cup \{v_1, \ldots, v_{\ell-1}\}$. We define v_ℓ to be u.
   (d) By Lemma 5.6, the total $\mathcal{W}_\ell$ weight of constraints that are incident on v_ℓ and only containing variables from $S^\star$ is at least $10 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$. We update h by setting $h(v_\ell)$ to be that value from $\{0,1\}$ such that at least half of the $\mathcal{W}_\ell$ weight of these constraints is satisfied.
3. Return the assignment h.

Figure 11: Procedure PERTURB for perturbing the optimal assignment

for MAX-w-SAT (but not for MAX-2-AND), is that it is possible to satisfy a constraint by just changing one of the variables it depends on. The partial assignment h is what we will be using to argue that Step 4. of the algorithm produces a good Pareto approximation. More formally, we have the following Lemma.

Lemma 5.7 *For the assignment h obtained from Procedure* PERTURB *(see Figure 11), for each $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$. Furthermore, for each* high variance *instance $\mathcal{W}_\ell$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant 4 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$.*

Proof: Consider the special variable v_ℓ that we choose for *high variance* instance $\ell \in [k]$. Since $v_\ell \notin B$, the constraints incident on v_ℓ only contribute at most a $\varepsilon/2k$ fraction of the objective value in each instance. Thus, changing the assignment v_ℓ can reduce the value of any instance by at most a $\frac{\varepsilon}{2k}$ fraction of their current objective value. Also, we pick different special variables for each *high variance* instance. Hence, the total effect of these perturbations on any instance is that it reduces the objective value (given by $h^\star \cup g$) by at most $1 - (1 - \frac{\varepsilon}{2k})^k \leqslant \frac{\varepsilon}{2}$ fraction. Hence for all instances $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$.

For a *high variance instance* $\ell \in [k]$, since $v_\ell \in U$, the variable v_ℓ must be one of the first $t/2$ variables brought into $S^\star$ because of ℓ. Hence, by Lemma 5.6 the total weight of constraints that are incident on v_ℓ and entirely contained inside $S^\star$ is at least $10 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$. Hence, there is an assignment to v_ℓ that satisfies at least at least half the weight of these MAX-w-SAT constraints[7]

[7]This is not true if they are MAX-2-AND constraints.

in ℓ. At the end of the iteration, when we pick an assignment to v_ℓ, we have $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant 5 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$. Since the later perturbations do not affect value of this instance by more than $\varepsilon/2$ fraction, we get that for the final assignment h, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot 5 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell) \geqslant 4 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$. ■

Given all this, we now show that with high probability the algorithm finds an assignment that satisfies, for each $\ell \in [k]$, at least $(3/4 - \varepsilon) \cdot c_\ell$ weight from instance $\mathcal{W}_\ell$. The following theorem immediately implies Theorem 1.3.

Theorem 5.8 *Let w be a constant. Suppose we're given $\varepsilon \in (0, 2/5]$, k simultaneous* Max-w-SAT *instances $\mathcal{W}_1, \ldots, \mathcal{W}_\ell$ on n variables, and target objective value $c_1, \ldots, c_k$ with the guarantee that there exists an assignment $f^\star$ such that for each $\ell \in [k]$, we have $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$. Then, the algorithm* Sim-MaxwSAT *runs in time $2^{O(k^3/\varepsilon^2 \log(k/\varepsilon^2))} \cdot \mathrm{poly}(n)$, and with probability at least 0.9, outputs an assignment f such that for each $\ell \in [k]$, we have, $\mathsf{val}(f, \mathcal{W}_\ell) \geqslant (3/4 - \varepsilon) \cdot c_\ell$.*

Proof: Consider the iteration of Step 4. of the algorithm when h_0 is taken to equal $h^\star$. Then, by Part 1 of Theorem 5.1, the LP in Step 4.a will be feasible (this uses the fact that $\mathsf{val}(f^\star, \mathcal{W}_\ell) \geqslant c_\ell$ for each ℓ).

By Lemma 5.5 and a union bound, with probability at least $1 - k\delta_0 > 0.9$, over the choice of g, we have that for *every* low variance instance $\ell \in [k]$, $\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) \geqslant (3/4 - \varepsilon/2) \cdot c_\ell$. Henceforth we assume that the assignment g sampled in Step 4.b of the algorithm is such that this event occurs. Let h be the output of the procedure Perturb given in Figure 11 for the input $h^\star$ and g. By Lemma 5.7, h satisfies

1. For every instance $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$.
2. For every high variance instance $\ell \in [k]$, $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant 4 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)$.

We now show that the desired Pareto approximation behavior is achieved when h is considered as the partial assignment in Step 4.c of the algorithm. We analyze the guarantee for low and high variance instances separately.

For any *low variance* instance $\ell \in [k]$, from property 1 above, we have $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (1 - \varepsilon/2) \cdot \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell)$. Since we know that $\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) \geqslant (3/4 - \varepsilon/2) \cdot c_\ell$, we have $\mathsf{val}(h \cup g, \mathcal{W}_\ell) \geqslant (3/4 - \varepsilon) \cdot c_\ell$.

For every high variance instance $\ell \in [k]$, since $h^\star = f^\star|_{S^\star}$, for any g we must have,

$$\mathsf{val}(h^\star \cup g, \mathcal{W}_\ell) \geqslant \mathsf{val}(f^\star, \mathcal{W}_\ell) - \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell) \geqslant c_\ell - \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell).$$

Combining this with properties 1 and 2 above, we get,

$$\begin{aligned} \mathsf{val}(h \cup g, \mathcal{W}_\ell) &\geqslant (1 - \varepsilon/2) \cdot \max\{c_\ell - \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell), 4 \cdot \mathsf{activedegree}_{S^\star}(\mathcal{W}_\ell)\} \\ &\geqslant (3/4 - \varepsilon) \cdot c_\ell. \end{aligned}$$

Thus, for all instances $\ell \in [k]$, we get $\mathsf{val}(h \cup g) \geqslant (3/4 - \varepsilon) \cdot c_\ell$. Since we are taking the best assignment $h \cup g$ at the end of the algorithm Sim-MaxwSAT, the theorem follows.

Running time : Running time of the algorithm is $2^{O(kt)} \cdot \mathrm{poly}(n)$ which is $2^{O(k^3/\varepsilon^2 \log(k/\varepsilon^2))} \cdot \mathrm{poly}(n)$. ■

References

[ABG06] Eric Angel, Evripidis Bampis, and Laurent Gourvs. Approximation algorithms for the bi-criteria weighted MAX-CUT problem. *Discrete Applied Mathematics*, 154(12):1685 – 1692, 2006.

[AGK+11] Noga Alon, Gregory Gutin, Eun Jung Kim, Stefan Szeider, and Anders Yeo. Solving MAX-r-SAT above a tight lower bound. *Algorithmica*, 61(3):638–655, 2011.

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.

[BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *FOCS*, pages 472–481, 2011.

[BS04] B. Bollobás and A. D. Scott. Judicious partitions of bounded-degree graphs. *Journal of Graph Theory*, 46(2):131–143, 2004.

[Cha13] Siu On Chan. Approximation resistance from pairwise independent subgroups. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing*, STOC '13, pages 447–456. ACM, 2013.

[CMM06] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Note on MAX-2SAT. *Electronic Colloquium on Computational Complexity (ECCC)*, 13(064), 2006.

[Dia11] Ilias Diakonikolas. *Approximation of Multiobjective Optimization Problems*. PhD thesis, Columbia University, 2011.

[DRS02] Irit Dinur, Oded Regev, and Clifford D. Smyth. The hardness of 3 - uniform hypergraph coloring. In *Proceedings of the 43rd Symposium on Foundations of Computer Science*, FOCS '02, pages 33–, Washington, DC, USA, 2002. IEEE Computer Society.

[GRW11] Christian Glaßer, Christian Reitwießner, and Maximilian Witek. Applications of discrepancy theory in multiobjective approximation. In *FSTTCS'11*, pages 55–65, 2011.

[GS11] Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with psd objectives. In *FOCS*, pages 482–491, 2011.

[GW93] Michel X. Goemans and David P. Williamson. A new $\frac{3}{4}$-approximation algorithm for MAX SAT. In *IPCO*, pages 313–321, 1993.

[GW95] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, November 1995.

[Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.

[IP01] Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-SAT. *Journal of Computer and System Sciences*, 62(2):367 – 375, 2001.

[IPZ01] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, December 2001.

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. pages 767–775, 2002.

[KO07] Daniela Kühn and Deryk Osthus. Maximizing several cuts simultaneously. *Comb. Probab. Comput.*, 16(2):277–283, March 2007.

[KSTW01] Sanjeev Khanna, Madhu Sudan, Luca Trevisan, and David P. Williamson. The approximability of constraint satisfaction problems. *SIAM J. Comput.*, 30(6):1863–1920, December 2001.

[Mar13] Dániel Marx. Slides : CSPs and fixed-parameter tractability. `http://www.cs.bme.hu/~dmarx/papers/marx-bergen-2013-csp.pdf`, 2013.

[MM12] Konstantin Makarychev and Yury Makarychev. Approximation algorithm for non-boolean max k-csp. In Anupam Gupta, Klaus Jansen, Jos Rolim, and Rocco Servedio, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, volume 7408 of *Lecture Notes in Computer Science*, pages 254–265. Springer Berlin Heidelberg, 2012.

[MR99] Meena Mahajan and Venkatesh Raman. Parameterizing above guaranteed values: Maxsat and maxcut. *J. Algorithms*, 31(2):335–354, 1999.

[MRS09] Meena Mahajan, Venkatesh Raman, and Somnath Sikdar. Parameterizing above or below guaranteed values. *J. Comput. Syst. Sci.*, 75(2):137–153, 2009.

[Pat08] Viresh Patel. Cutting two graphs simultaneously. *J. Graph Theory*, 57(1):19–32, January 2008.

[PY00] Christos H. Papadimitriou and Mihalis Yannakakis. On the approximability of tradeoffs and optimal access of web sources. In *Foundations of Computer Science, 2000. Proceedings. 41st Annual Symposium on*, pages 86–92, 2000.

[Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 245–254, New York, NY, USA, 2008. ACM.

[RS04] Dieter Rautenbach and Zoltán Szigeti. *Simultaneous large cuts*. Forschungsinstitut für Diskrete Mathematik, Rheinische Friedrich-Wilhelms-Universität, 2004.

[RS09] Prasad Raghavendra and David Steurer. How to round any CSP. In *In Proc. 50th IEEE Symp. on Foundations of Comp. Sci*, 2009.

[RT12] Prasad Raghavendra and Ning Tan. Approximating csps with global cardinality constraints using sdp hierarchies. In *SODA*, pages 373–387, 2012.

[Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC '78, pages 216–226, New York, NY, USA, 1978. ACM.

[Tre98] L. Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998.

A Hardness results for large k

In this section, we prove our hardness results for simultaneous CSPs. Recall the theorem that we are trying to show.

Theorem A.1 (restated) *Assume the Exponential Time Hypothesis. Let $\mathcal{F}$ be a fixed finite set of Boolean predicates. If $\mathcal{F}$ is not 0-valid or 1-valid, then for $k = \omega(\log n)$, then detecting positivity of k-fold simultaneous MAX-$\mathcal{F}$-CSPs on n variables requires time superpolynomial in n.*

The main notion that we will use for our hardness reductions is the notion of a "simultaneous-implementation".

Definition A.2 (Simultaneous-Implementation) *Let $\{x_1, \ldots, x_w\}$ be a collection of variables (called* **primary** *variables). Let $P : \{0,1\}^w \to \{\text{TRUE}, \text{FALSE}\}$ be a predicate. Let $\{y_1, \ldots, y_t\}$ be another collection of variables (called* **auxiliary** *variables).*

Let $\mathcal{C}_1, \ldots, \mathcal{C}_k$ be sets of constraints on $\{x_1, \ldots, x_w, y_1, \ldots, y_t\}$, where for each $i \in [k]$, $\mathcal{C}_i$ consists of various applications of predicates to tuples of distinct variables from $\{x_1, \ldots, x_w, y_1, \ldots, y_t\}$. We say that $\mathcal{C}_1, \ldots, \mathcal{C}_k$ **simultaneously-implements** *P if for every assignment to $x_1, \ldots, x_w$, we have,*

- *If $P(x_1, \ldots, x_w) = \text{TRUE}$, then there exists a setting of the variables $y_1, \ldots, y_t$ such that each collection $\mathcal{C}_1, \ldots, \mathcal{C}_k$ has at least one satisfied constraint.*
- *If $P(x_1, \ldots, x_w) = \text{FALSE}$, then for every setting of the variables $y_1, \ldots, y_t$, at least one of the collections $\mathcal{C}_1, \ldots, \mathcal{C}_k$ has no satisfied constraints.*

We say that a collection of predicates $\mathcal{F}$ simultaneously-implements P if there is a simultaneous-implementation of P where for each collection $\mathcal{C}_i$ ($i \in [k]$), every constraint in $\mathcal{C}_i$ is an application of some predicate from $\mathcal{F}$.

The utility of simultaneous-implementation lies in the following lemma.

Lemma A.3 *Let P be a predicate. Suppose checking satisfiability of CSPs on n variables with m constraints, where each constraint is an application of the predicate P, requires time $T(n, m)$, with $T(n, m) = \omega(m + n)$. Suppose $\mathcal{F}$ simultaneously-implements P. Then detecting positivity of $O(m)$-fold simultaneous MAX-$\mathcal{F}$-CSP on $O(m + n)$ variables requires time $\Omega(T(n, m))$.*

Proof: Suppose we have a P-CSP instance Φ with m constraints on n variables. For each of the constraints $C \in \Phi$, we simultaneously-implement C using the original set of variables as primary variables, and new auxiliary variables for each constraint. Thus, for every $C \in \Phi$, we obtain k MAX-$\mathcal{F}$-CSP instances $\mathcal{C}_1^C, \ldots, \mathcal{C}_k^C$, for some constant k. The collection of instances $\{\mathcal{C}_i^C\}_{C \in \Phi, i \in [k]}$ constitute the $O(m)$-simultaneous MAX-$\mathcal{F}$-CSP instance on $O(m + n)$ variables.

If Φ is satisfiable, we know that there exists an assignment to the original variables such that each $C \in \Phi$ is satisfied. Hence, by the simultaneously-implements property, there exists as assignment to all the auxiliary variables such that each $\mathcal{C}_i^C$ has at least one satisfied constraint. If Φ is unsatisfiable, for any assignment to the primary variables, at least one constraint C must be unsatisfied. Hence, by the simultaneously-implements property, for any assignment to the auxiliary variables, there is an $i \in [k]$ such that $\mathcal{C}_i^C$ has no satisfied constraints. Thus, our simultaneous MAX-$\mathcal{F}$-CSP instance has a non-zero objective value iff Φ is satisfiable. Since this reduction requires only $O(m+n)$ time, suppose we require T' time for detecting positivity of a $O(m)$-simultaneous MAX-$\mathcal{F}$-CSP instance on $O(m + n)$ variables, we must have $T' + O(m + n) \geqslant T(m, n)$, giving $T' = \Omega(T(m, n))$ since $T(m, n) = \omega(m + n)$. ■

The simultaneous-implementations we construct will be based on a related notion of implementation arising in approximation preserving reductions. We recall this definition below.

Definition A.4 (Implementation) *Let $x_1, \ldots, x_w$ be a collection of variables (called* **primary** *variables). Let $P : \{0,1\}^w \to \{\text{TRUE}, \text{FALSE}\}$ be a predicate.*

Let $y_1, \ldots, y_t$ be another collection of variables (called **auxiliary** *variables). Let $C_1, \ldots, C_d$ be constraints on $\{x_1, \ldots, x_w, y_1, \ldots, y_t\}$, where for each $i \in [d]$, the variables feeding into C_i are all distinct.*

*We say that $C_1, \ldots, C_d$ e-**implements** P if for every assignment to $x_1, \ldots, x_w$ we have,*

- *If $P(x_1, \ldots, x_w) = \text{TRUE}$, then there exists a setting of the variables $y_1, \ldots, y_t$ such that at least e of the constraints $C_1, \ldots, C_d$ evaluate to* TRUE.
- *If $P(x_1, \ldots, x_w) = \text{FALSE}$, then for every setting of the variables $y_1, \ldots, y_t$, at most $e - 1$ of the constraints $C_1, \ldots, C_d$ evaluate to* TRUE.

We say that a collection of predicates $\mathcal{F}$ implements P if there is some e and an e-implementation of C where all the constraints $C_1, \ldots, C_d$ come from $\mathcal{F}$.

We will be using following predicates in our proofs.

- $\mathsf{Id}, \mathsf{Neg}$: These are the unary predicates defined as $\mathsf{Id}(x) = x$ and $\mathsf{Neg}(x) = \bar{x}$.
- NAE: w-ary NAE predicate on variables $x_1, \ldots, x_w$ is defined as $\mathsf{NAE}(x_1, \ldots, x_w) = \text{FALSE}$ iff all the x_i's are equal.
- $\mathsf{Equality}$: $\mathsf{Equality}$ is a binary predicate given as $\mathsf{Equality}(x, y) = \text{TRUE}$ iff x equals y.

We will use the following Lemmas from [KSTW01].

Lemma A.5 ([KSTW01]) *Let f be a predicate which is not 0-valid, and which is closed under complementation. Then $\{f\}$ implements $\mathsf{XOR}(x, y)$.*

Lemma A.6 ([KSTW01]) *Let f be a predicate not closed under complementation, and let g be a predicate that is not 0-valid. Then $\{f, g\}$ implements Id, and $\{f, g\}$ implements Neg.*

We will now prove lemmas that will capture the property of *simultaneous implementation* which will be used in proving Theorem A.1.

Lemma A.7 *If $\{f\}$ simultaneously-implements predicate XOR on two variables, then $\{f\}$ also simultaneously-implements the predicate NAE on three variables.*

Proof: Consider an NAE constraint $\mathsf{NAE}(x, y, z)$. Let $\mathcal{A}_1, \ldots, \mathcal{A}_d$ be the simultaneous implementation of constraint $\mathsf{XOR}(x, y)$, using predicate f and a set of auxiliary variables $y_1, \ldots, y_t$ for some t. Similarly, let $\mathcal{B}_1, \ldots, \mathcal{B}_d$ and $\mathcal{C}_1, \ldots, \mathcal{C}_d$ be the simultaneous implementation of constraint $\mathsf{XOR}(y, z)$ and $\mathsf{XOR}(x, z)$ respectively using f and on a same set of auxiliary variables $y_1, \ldots, y_t$, constructed by replacing the variables (x, y) in $\{\mathcal{A}_1, \ldots, \mathcal{A}_d\}$ with (y, z) and (x, z) respectively. We construct sets of constraints $\mathcal{D}_1, \ldots, \mathcal{D}_d$ as follows: for each $i \in [d]$, $\mathcal{D}_i$ consists of all constraints from $\mathcal{A}_i, \mathcal{B}_i$, and $\mathcal{C}_i$. We now show that $\{\mathcal{D}_1, \ldots, \mathcal{D}_d\}$ simultaneously-implement $\mathsf{NAE}(x, y, z)$.

First, notice that $\mathsf{NAE}(x, y, z)$ is FALSE iff all constraints $\mathsf{XOR}(x, y)$, $\mathsf{XOR}(y, z)$ and $\mathsf{XOR}(x, z)$ are FALSE. Consider the case when $\mathsf{NAE}(x, y, z)$ is FALSE. Since we are using same set of auxiliary

variables and the implementation is symmetric, for every setting of variables $y_1, \ldots, y_t$, there exists a fixed $i \in [d]$ such that each of $\mathcal{A}_i, \mathcal{B}_i$ and $\mathcal{C}_i$ has no satisfied constraints. And hence, instance $\mathcal{D}_i$ has no satisfied constraints. If $\mathsf{NAE}(x, y, z)$ is TRUE then at least one of $\mathsf{XOR}(x, y)$, $\mathsf{XOR}(y, z)$ or $\mathsf{XOR}(x, z)$ must be TRUE. Without loss of generality, we assume that $\mathsf{XOR}(x, y)$ is TRUE. Thus, there exists a setting of variables $y_1, \ldots, y_t$ such that each of $\mathcal{A}_1, \ldots, \mathcal{A}_d$, has at least one satisfied constraint, and hence each of $\mathcal{D}_1, \ldots, \mathcal{D}_d$ too has at least one such constraint. ■

Lemma A.8 *Let f be a predicate not closed under complementation, not* 0*-valid and not* 1*-valid. f can simultaneously-implement* Equality.

Proof: Consider an equality constraint $\mathsf{Equality}(x, y)$, our aim is to simultaneously-implement this constraint using predicate f.

Since f satisfies the properties of Lemma A.6, we can implement $\mathsf{Id}(x)$ and $\mathsf{Id}(y)$ using f. Let $X_1^T, \ldots, X_{d_1}^T$ be an e_1-implementation of $\mathsf{Id}(x)$ using f and some set of auxiliary variables A_1 for some $e_1 < d_1$. Similarly, let $Y_1^T, \ldots, Y_{d_1}^T$ be an e_1-implementation of $\mathsf{Id}(y)$ using f and a set of auxiliary variables A_2.

We can also implement $\mathsf{Neg}(x)$ and $\mathsf{Neg}(y)$ using f. Let $X_1^F, \ldots, X_{d_2}^F$ be an e_2-implementation of $\mathsf{Neg}(x)$ using f and a set of auxiliary variables B_1 for some $e_2 < d_2$. Similarly, let $Y_1^F, \ldots, Y_{d_1}^F$ be an e_2-implementation of $\mathsf{Neg}(y)$ using f and a set of auxiliary variables B_2.

We now describe the construction of the simultaneous-implementation. The implementation uses all auxiliary variables in A_1, A_2, B_1, and B_2. Each instance in the simultaneous-implementation is labeled by a tuple (M, N, a, b) where $M \subseteq [d_1]$ with $|M| = d_1 - e_1 + 1$, $N \subseteq [d_2]$ with $|N| = d_2 - e_2 + 1$, and $(a, b) \in \{(T, F), (F, T)\}$. An instance corresponding to a tuple (M, N, a, b) has following set of constraints in f:

$$\{X_m^a, Y_m^b | m \in M, n \in N\}$$

We will now prove the simultaneous-implementation property of the above created instance. Consider the case when $x = y =$ TRUE (other case being similar). We know that in this case, there exists a setting of auxiliary variables A_1 used in the implementation of $\mathsf{Id}(x)$ which satisfies at least e_1 constraints out of $X_1^T, \ldots, X_{d_1}^T$. Similarly, there exists a setting of auxiliary variables A_2 used in the implementation of $\mathsf{Id}(y)$ which satisfies at least e_1 constraints out of $Y_1^T, \ldots, Y_{d_1}^T$. Fix this setting of auxiliary variables in A_1, A_2, and any arbitrary setting for auxiliary variables in B_1 and B_2. Thus, the instance labeled by tuple the (M, N, a, b) either contains $d_1 - e_1 + 1$ constraints from $X_1^T, \ldots, X_{d_1}^T$ if $a = T$, or else, it contains $d_1 - e_1 + 1$ constraints from $Y_1^T, \ldots, Y_{d_1}^T$. In any case, the property of e_1-implementation implies that at least one constraint is satisfied for this instance.

Now we need to show that if $x \neq y$, then for any setting of auxiliary variables, there exists an instance which has no satisfied constraints. Consider the case when $x =$ TRUE and $y =$ FALSE (other case being similar). Consider any fixed assignment to the auxiliary variables in A_1, A_2, B_1, and B_2. We know that for this fixed assignment to the auxiliary variables in B_1, there exists a subset $N \subseteq [d_2]$ of size at least $d_2 - e_2 + 1$, such that all constraints in $\{X_j^F | j \in N\}$ are unsatisfied. Similarly, for this fixed assignment to variables in A_2, there exists a subset $M \subseteq [d_1]$ of size at least $d_1 - e_1 + 1$ such that all constraints in $\{Y_i^T | i \in M\}$ are unsatisfied. Thus, the instance corresponding to tuple (M, N, F, T) has no satisfied constraints. ■

We now prove Theorem A.1.

Proof: We take cases on whether $\mathcal{F}$ contains some f which is closed under complementation.

Case 1: Suppose there exists some $f \in \mathcal{F}$ which is closed under complementation. In this case, it is enough to show that f simultaneously-implements XOR. To see this, assume that we can simultaneously-implement XOR using f. Hence, by Lemma A.7, we can simultaneously-implement the predicate NAE on *three* variables using f. We start with an NAE-3-SAT instance ϕ, on n variables with m constraints. For each constraint $C \in \phi$, we create a set of $O(1)$ many instances which simultaneously-implement C. The final simultaneous instance is the collection of all instances that we get with each simultaneous-implementation of constraints in ϕ.

In the completeness case, when ϕ is satisfiable, then by the property of simultaneous-implementation, we have that there exists a setting of auxiliary variables, from each implementation of NAE constraints, such that each instance has at least one constraint satisfied. And hence, the value of the final simultaneous instance is non *zero*.

In the soundness case, for any assignment to the variables $x_1, \ldots, x_n$ there exists a constraint (say C) which is not satisfied. Hence one of the instance from the simultaneous implementation of this constraint has value *zero* no matter how we set the auxiliary variables. And hence, the whole simultaneous instance has value zero in this case.

To prove the theorem in this case, it remains to show that we can simultaneously-implement $\mathsf{XOR}(x, y)$ using f. Since f is closed under complementation, we can e-implement XOR using f (for some e) by Lemma A.5. Let $C_1, \ldots, C_d$ be the set of f-constraints that we get from this e-implementation, $e < d$. The collection of instances contains one instance for every subset $J \subseteq [d]$ of size $d - e + 1$. The instance labeled by $J \subseteq [d]$ contains all constraints from the set $\{C_j | j \in J\}$. Hence, there $\binom{d}{e-1}$ instances in the collection. Note that we used the same set of auxiliary variables in this simultaneous-implementation. We now show that this collection of instances simultaneously-implements $\mathsf{XOR}(x, y)$. To see this, consider the case when $\mathsf{XOR}(x, y)$ is TRUE. Thus. there is an assignment to the auxiliary variables that satisfies at least e constraints out of $C_1, \ldots, C_d$. Hence, for this particular assignment, the instance labeled by J, where $J \subset [d]$ is any subset of size $d-e+1$, has at least one satisfied constraint. When $\mathsf{XOR}(x, y)$ is FALSE, then for any assignment to the auxiliary variables, there is some $J \subseteq [d]$ of size $d - e + 1$ such that no constraints in the set $\{C_j | j \in J\}$ are satisfied. Hence, for this assignment, the instance labeled with J has no satisfied constraints. This shows that f simultaneously-implements predicate XOR on two variables.

Combining the two arguments above, we get that $\{f\}$ simultaneously-implements 3-NAE. Since 3-NAE has a linear time gadget reduction from 3-SAT [Sch78], and the ETH implies that 3-SAT on s variables and $O(s)$ clauses requires time $2^{\Omega(s)}$ [IP01, IPZ01], we get that checking satisfiability of a 3-NAE instance with $\omega(\log n)$ constraints on $\omega(\log n)$ variables requires time super-polynomial in n. Thus, using Lemma A.3 implies that detecting positivity of an $\omega(\log n)$-simultaneous MAX-f-CSP requires time superpolynomial in n.

Case 2: Suppose that for all $f \in \mathcal{F}$, f is not closed under complementation. Let $f \in \mathcal{F}$ be any predicate of arity r. Since, f is not closed under complementation, there exist $\alpha, \beta \in \{0,1\}^r$ that satisfy $\alpha_i \oplus \beta_i = 1$ for all $i \in [r]$, and $f(\alpha) = 0$, $f(\beta) = 1$. We can reduce a 3-SAT instance with n variables and $m = \text{poly}(n)$ clauses to m simultaneous instances over n variables involving the predicate f. For every clause C of the form $x \vee y \vee z$, we create an instance with 3 equal weight constraints $\{f(\alpha \oplus (x, \ldots, x)), f(\alpha \oplus (y, \ldots, y)), f(\alpha \oplus (z, \ldots, z))\}$, where $\oplus$ denotes bitwise-xor, or equivalently, we negate the variable in the i-th position iff $\alpha_i = 1$.

It is straightforward to see that the original 3-SAT formula is satisfiable if and only if there is an assignment to the variables that simultaneously satisfies a non zero fraction of the constraints

in each of the instances.

In the above reduction, we must be able to apply the predicate to several copies of the same variable. In order to remove this restriction, we replace each instance with a collection $\mathcal{C}$ of instances as follows: Consider an instance $\{f(\alpha \oplus (x, \ldots, x)), f(\alpha \oplus (y, \ldots, y)), f(\alpha \oplus (z, \ldots, z))\}$. We add to our collection $\mathcal{C}$, an instance $\{f(\alpha \oplus (a_1, \ldots, a_r)), f(\alpha \oplus (b_1, \ldots, b_r)), f(\alpha \oplus (c_1, \ldots, c_r))\}$, where a_i, b_i and c_i for all $i \in [r]$, are the fresh set of variables. Using Lemma A.8, we can simultaneously-implement each constraint of the form $x = a_i$, $y = b_i$ and $z = c_i$ using f. We add all the instances obtained from the simultaneous-implementations to the collection $\mathcal{C}$. Notice that, we have replaced each original instance with only $O(1)$ many instances. Hence, we have $O(m)$ many instances in our final construction. Thus, as in the first case, assuming ETH we deduce that detecting positivity of an $\omega(\log n)$-simultaneous MAX-f-CSP requires time super-polynomial in n. ■

A.1 Hardness for Simultaneous Max-w-SAT

Proposition A.9 (Proposition 1.1 restated) *For all integers $w \geqslant 4$ and $\varepsilon > 0$, given $k \geqslant 2^{w-3}$ simultaneous instances of* Max-Ew-SAT *that are simultaneously satisfiable, it is* **NP***-hard to find a $(7/8 + \varepsilon)$-minimium approximation.*

Proof: We know that given a satisfiable Max-E3-SAT instance, it is **NP**-hard to find an assignment that satisfies a $(7/8+\varepsilon)$ fraction of the constraints [Hås01]. We reduce a single Max-E3-SAT instance to the given problem as follows : Let Φ be an instance of Max-E3-SAT with clauses $\{C_i\}_{i=1}^m$ on variable set $\{x_1, \ldots, x_n\}$. Given $w \geqslant 4$, let $\{z_1, \ldots, z_{w-3}\}$ be a fresh set of variables. For every, $a \in \{0,1\}^{w-3}$, we construct a Max-Ew-SAT instance with clauses $\{C_i \vee \vee_{j=1}^{w}(z_j \oplus a_j)\}_{i=1}^m$, where $z_j \oplus 0 = z_j$ and $z_j \oplus 1 = \bar{z_j}$. It is straightforward to see that for any assignment, its value on Φ is the same as the minimum of its value on the Max-Ew-SAT instances, immediately implying the result. ■

B Algorithm for Unweighted Max-CUT

For simultaneous unweighted Max-CUT instances, we can use the Goemans-Williamson SDP to obtain a slightly better approximation. The algorithm, UnweightedMC, is described in Figure 12.

Let V be the set of vertices. Our input consists of an integer $k \geqslant 1$, and k unweighted instances of Max-CUT, specified by indicator functions $\mathcal{W}_1, \ldots, \mathcal{W}_k$ of edge set. Let m_ℓ denotes the number of edges in graph $\ell \in [k]$. We consider these graphs as weighted graphs with all non-zero edge weights as $\frac{1}{m_\ell}$ so that the total weight of edges of in a graph is 1. For a given subset S of vertices, we say an edge is *active* if at least one of its endpoints is in $V \backslash S$.

B.1 Analysis of Sim-UnweightedMC

For analysing the algorithm Sim-UnweightedMC, we need the following lemma that is proven by combining SDP rounding for 2-SAT from [CMM06] with a Markov argument. A proof is included in Section C for completeness.

Lemma B.1 *For k simultaneous instances of any* MAX-2-CSP *such that there exists an assignment which satisfies a $1 - \varepsilon$ weight of the constraints in each of the instances, there is an efficient algorithm that, for n large enough, given an optimal partial assignment h to a subset of variables,*

Input: k unweighted instances of MAX-CUT $\mathcal{W}_1, \ldots, \mathcal{W}_k$ on the vertex set V.
Output: A cut of V.

1. Set $\varepsilon \stackrel{\text{def}}{=} \frac{1}{1600 \cdot c_0^2 k^2}, t = \frac{100k}{\varepsilon^2}, S = \varnothing, D = \varnothing$ (c_0 is the constant from Lemma B.1).
2. If every graph has more than t edges, then go to Step 4..
3. Repeat until there is no $\ell \in [k] \backslash D$ such that the instance $\mathcal{W}_\ell$ has less than $t^{3^{|D|}}$ active edges given S.
   (a) Let $\ell \in [k]$ be an instance with the least number of edges active edges given S.
   (b) Add all the endpoints of the edge set of instance $\mathcal{W}_\ell$ into set S.
   (c) $D \leftarrow D \cup \ell$
4. For each partial assignment $h : S \to \{0,1\}$ (If $S = \varnothing$ then do the following steps without considering partial assignment h)
   (a) Run the SDP algorithm for instances in $[k]$ given by Lemma B.1 with h as a partial assignment. Let h_1 be the assignment returned by the algorithm. (Note $h_1|_S = h$)
   (b) Define $g : V \backslash S \to \{0,1\}$ by independently sampling $g(v) \in \{0,1\}$ with $\mathbf{E}[g(v)] = 1/2$, for each $v \in V \backslash S$. In this case the cut is given by an assignment $h \cup g$.
   (c) Let out_h be the better of the two solutions (h_1 and $h \cup g$).
5. Output the largest out_h seen.

Figure 12: Algorithm SIM-UNWEIGHTEDMC for approximating unweighted simultaneous MAX-CUT

returns a full assignment which is consistent with h and simultaneously satisfies at least $1 - c_0 k\sqrt{\varepsilon}$ (for an absoute constant c_0) fraction of the constraints in each instance with probability 0.9.

Let $S^\star, D^\star$ denote the set S and D that we get at the end of step 3 of the algorithm SIM-UNWEIGHTEDMC. Let $f^\star : V \to \{0,1\}$ denote the optimal assignment and let $h^\star = f^\star|_{S^\star}$.

Theorem B.2 *For large enough n, given k simultaneous unweighted* MAX-CUT *instances on n vertices, the algorithm* SIM-UNWEIGHTEDMC *returns computes a $\left(\frac{1}{2} + \Omega\left(\frac{1}{k^2}\right)\right)$-minimum approximate solution with probability at least* 0.9. *The running time is $2^{2^{2^{O(k)}}} \cdot \mathrm{poly}(n)$.*

Proof: We will analyze the approximation guarantee of the algorithm when the optimal partial assignment $h^\star$ to the variables $S^\star$ is picked for h in Step 4. of the algorithm. Note that Step 4.a and 4.b maintain the assignment to the set $S^\star$ given in Step 4. Hence, for all instances $\ell \in D^\star$, we essentially get the optimal cut value $\mathsf{val}(f^\star, \mathcal{W}_\ell)$. We will analyze the effect of rounding done in Step 4.a and 4.b on instances in $[k] \backslash D^\star$ for a partial assignment $h^\star$ to $S^\star$. Since we are taking the best of the two roundings, it is enough to show the claimed guarantee for at least one of these two steps.

Let OPT be the value of optimal solution for a given set of instances $[k]$. We consider two cases depending on the value of this optimal solution.

1. OPT $\geqslant (1-\varepsilon)$: In this case, we show that the cut returned in Step 4.a is good with high probability.

Since the OPT is at least $(1-\varepsilon)$, and $h^\star$ is an optimal partial assignment, we can apply Lemma B.1 such that with probability at least 0.9 we get a cut of value at least $(1-10c_0k\cdot\sqrt{\varepsilon})$ for all graphs $\ell \in [k]\backslash D^\star$, for some constant c_0. In this case, the approximation guarantee is at least :

$$(1-10c_0k\cdot\sqrt{\varepsilon}) \geqslant \frac{3}{4}.$$

2. OPT $< (1-\varepsilon)$: In this case, we show that the cut returned in Step 4.b gives the claimed approximation guarantee with high probability.

Fix a graph $\ell \in [k]\backslash D^\star$, if any. Let m_ℓ be the number of edges in this graph. We know that $m_\ell \geqslant t^{3^{|D^\star|}}$ and also $|S^\star| \leqslant 4t^{3^{|D^\star|-1}}$. Let Y_ℓ be a random variable defined as

$$Y_\ell \stackrel{\text{def}}{=} \mathsf{val}(h^\star \cup g, \mathcal{W}_\ell),$$

that specifies the fraction of *total* edges that are cut by assignment $h^\star \cup g$ where g is a random partition g of a vertex set $V\backslash S^\star$. The number of edges of graph ℓ that are not active given $S^\star$ is at most $1/2\cdot|S^\star|^2$. If $|D^\star|=0$, we know that all the edges in graph ℓ are active. Otherwise, using the bounds on m_ℓ and $|S^\star|$, we get that at least a $(1-1/t)$ fraction of the total edges are active given $S^\star$. This implies that for uniformly random partition g,

$$\mathbf{E}_g[Y_\ell] \geqslant \frac{1}{2}\sum_{\substack{C\in\mathcal{C}\\ C\in\mathsf{Active}(S^\star)}} \mathcal{W}_\ell(C) \geqslant 1/2(1-1/t).$$

We now analyze the variance of a random variable Y_ℓ under uniformly random assignment $g : V\backslash S^\star \to \{0,1\}$.

$$\mathbf{Var}_g[Y_\ell] = \sum_{C_1,C_2\in\mathsf{Active}(S^\star)} \mathcal{W}(C_1)\mathcal{W}(C_2)\cdot(\mathbf{E}[C_1(h^\star\cup g)C_2(h^\star\cup g)] - \mathbf{E}[C_1(h^\star\cup g)]\mathbf{E}[C_2(h^\star\cup g)]).$$

The term in the above summation is *zero* unless we have either $C_1 = C_2$ (in which case we know $\mathbf{E}[C_1(h^\star\cup g)C_2(h^\star\cup g)] - \mathbf{E}[C_1(h^\star\cup g)]\mathbf{E}[C_2(h^\star\cup g)] = 1/4$) or when the edges C_1 and C_2 have a common endpoint in $V\backslash S^\star$ and the other endpoint in $S^\star$ (in this case $\mathbf{E}[C_1(h^\star\cup g)C_2(h^\star\cup g)] - \mathbf{E}[C_1(h^\star\cup g)]\mathbf{E}[C_2(h^\star\cup g)] \leqslant 1/4$). For $v\in V\backslash S^\star$, let κ_v be the set of edges whose one endpoint is v and other endpoint in $S^\star$. Thus,

$$\begin{aligned}
\mathbf{Var}_g[Y_\ell] &\leqslant \frac{1}{4}\sum_{C\in\mathsf{Active}(S^\star)} \mathcal{W}(C)^2 + \frac{1}{4}\sum_{v\in V\backslash S^\star}\sum_{C_1,C_2\in\kappa_v} \mathcal{W}(C_1)\mathcal{W}(C_2)\\
&= \frac{1}{4m_\ell} + \frac{1}{4}\frac{1}{m_\ell^2}\sum_{v\in V\backslash S^\star}|\kappa_v|^2\\
&\leqslant \frac{1}{4m_\ell} + \max_{v\in V\backslash S^\star}|\kappa_v|\cdot\frac{1}{4}\frac{1}{m_\ell^2}\sum_{v\in V\backslash S^\star}|\kappa_v|\\
&\leqslant \frac{1}{4m_\ell} + |S^\star|\cdot\frac{1}{4}\frac{1}{m_\ell^2}\cdot m_\ell
\end{aligned}$$

$$\leqslant \frac{1}{4m_\ell} + \frac{1}{4}\frac{|S^\star|}{m_\ell}$$
$$\leqslant \frac{1}{4t^{3^{|D^\star|}}} + \frac{1}{4}\frac{|S^\star|}{t^{3^{|D^\star|}}} \leqslant \frac{1}{2t}.$$

Hence, by Chebyshev's Inequality, we have

$$\Pr\left[Y_\ell < \frac{1}{2}\cdot(1-\varepsilon_0-1/t)\right] \leqslant \frac{4\,\mathbf{Var}_g[Y_\ell]}{\varepsilon_0^2} \leqslant \frac{4\cdot 1/2t}{\varepsilon_0^2} \leqslant \frac{2}{\varepsilon_0^2 t}.$$

By a union bound, with probability at least $1-\frac{2k}{\varepsilon_0^2\cdot t}$, we get a simultaneous cut of value at least $\frac{1}{2}\cdot(1-\varepsilon_0-1/t)$ for all $\ell\in[k]\backslash D^\star$. If we take $\varepsilon_0=\frac{\sqrt{20k}}{\sqrt{t}}$, then with probability at least 0.9 we get a cut of value at least $\frac{1}{2}\cdot(1-\varepsilon_0-1/t)$ for all $\ell\in[k]\backslash D^\star$. In this case, the approximation guarantee is at least

$$\frac{\frac{1}{2}\cdot(1-\varepsilon_0-1/t)}{\left(1-\frac{1}{(40c_0k)^2}\right)} = \left(\frac{1}{2}+\Omega\left(\frac{1}{k^2}\right)\right).$$

■

C Semidefinite Programs for Simultaneous Instances

In this section, we study Semidefinite Programming (SDP) relaxations for simultaneous MAX-2-CSP instances.

C.1 Integrality gaps for Simultaneous MAX-CUT SDP

In this section, we show the integrality gaps associated with the natural SDP of *minimum approximation* problem for k-fold simultaneous MAX-CUT.

Suppose we have k simultaneous MAX-CUT instances on the set of vertices $V=\{x_1,\ldots,x_n\}$, specified by the associated weight functions $\mathcal{W}_1,\ldots,\mathcal{W}_k$. As before, let $\mathcal{C}$ denotes the set of all possible edges on V. We assume that for each $\ell\in[k]$, $\sum_{C\in\mathcal{C}}\mathcal{W}_\ell(C)=1$. Following Goemans and Williamson [GW95], the semi-definite programming relaxation for such an instance is described in Figure 13. We now prove the following claims about integrality gap for the above SDP.

$$\begin{array}{lrcll} & \text{maximize} & t & & \\ \text{s.t.} \quad \displaystyle\sum_{\substack{C\in\mathcal{C}\\ C=(x_i,x_j)}} \frac{1}{2}\cdot\mathcal{W}_\ell(C)\cdot(1-\langle v_i,v_j\rangle) & \geqslant & t & & \forall \ell\in[k] \\ \|v_i\|^2 & = & 1 & \text{for } i=1,\ldots,n & \end{array}$$

Figure 13: Semidefinite Program (SDP) for minimum approximation Simultaneous MAX-CUT

Claim C.1 *For weighted instances, the SDP for minimum approximation of simultaneous* MAX-CUT *does not have* any *constant integrality gap.*

Proof: Consider 3 simultaneous instances such that all but a tiny fraction of the weight of instance i is on edge i of a 3-cycle. Clearly, no cut can simultaneously cut all the three edges in the three cycle, and hence the optimum is tiny. However, for the simultaneous SDP, a vector solution that assigns to the three vertices of the cycle three vectors such that $\langle v_i, v_j \rangle = -1/2$ for $i \neq j$ gives a constant objective value for all three instances. ■

Claim C.2 *For every fixed k, there exists k-instances of* Max-CUT *where the SDP relaxation has value $1 - \Omega\left(\frac{1}{k^2}\right)$, while the maximum simultaneous cut has value only $\frac{1}{2}$. Moreover, the random hyperplane rounding for a good vector solution for this instance, returns a simultaneous cut of value 0.*

Proof: Let k be odd. We define k graphs on kn vertices. Partition the vertex set into $S_0, S_1, ..., S_{k-1}$, each of size n. Graph G_i has only edges (x, y) such that $x \in S_i$ ans $y \in S_{(i+1) \mod k}$, each of weight $1/n^2$. The optimal cut must contain exactly half the number of vertices from each partition, giving a simultaneous cut value of $1/2$. Whereas, the following SDP vectors achieve a simultaneous objective of $\left(1 - O(\frac{1}{k^2})\right)$: For all vertices in S_i, we assign the vector $\left(\cos \frac{i}{k}\pi, \sin \frac{i}{k}\pi\right)$. It is straightforward to see that applying the hyperplane rounding algorithm to this vector solution gives (with probability 1) a simultaneous cut value of 0. ■

C.2 SDP for Simultaneous Max-CSP

For Max-CSP, we will be interested in the regime where the optimum assignment satisfies at least a $(1 - \varepsilon)$ fraction of the constraints in each of the instances.

Given a MAX-2-CSP instance, we use the standard reduction to transform it into a MAX-2-SAT instance: We reduce each constraint of the 2-CSP instance with a set of at most 4 2-SAT constraints such that for any fixed assignment, the 2-CSP constraint is satisfied iff all the 2-SAT constraints are satisfied, and if the 2-CSP constraint is not satisfied, then at least one of the 2-SAT constraint is not satisfied. *e.g.* We replace $x_1 \wedge x_2$ with $x1 \vee x_2, \overline{x_1} \vee x_2$, and $x_1 \vee \overline{x_2}$. Similarly, we replace $x_1 \neq x_2$ with $x_1 \vee x_2$ and $\overline{x_1} \vee \overline{x_2}$. We distribute the weight of the 2-CSP constraint equally amongst the 2-SAT constraints.

Given k simultaneous MAX-2-CSP instances, we apply the above reduction to each of the instances to obtain k simultaneous Max-2-SAT instances. The above transformation guarantees the following:

- **Completeness** If there was an assignment of variables that simultaneously satisfied all the constraints in each of the MAX-2-CSP instances, then the same assignment satisfies all the constraints in each of the Max-2-SAT instances.
- **Soundness** If no assignment of variables simultaneously satisfied more than $(1 - \varepsilon)$ weighted fraction of the constraints in each of the MAX-2-CSP instances, then no assignment simultaneously satisfies more than $(1 - \varepsilon/4)$ weighted fraction of the constraints in each of the MAX-2-SAT instances.

From now on, we will assume that we have k simultaneous Max-2-SAT instances on the set of variables $\{x_1, \ldots, x_n\}$, specified by the associated weight functions $\mathcal{W}_1, \ldots, \mathcal{W}_k$. As before $\mathcal{C}$ denotes the set of all possible 2-SAT constraints on V. We assume that for each $\ell \in [k]$, $\sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) = 1$. Following Charikar *et al.* [CMM06], the semi-definite programming relaxation for such an instance is described in Figure 14.

For convenience, we replace each negation $\bar{x_i}$ with a new variable x_{-i}, that is equal to $\overline{x_1}$ by definition. For each variable $x_i \in V$, the SDP relaxation will have a vector v_i. We define $v_{-i} = -v_i$. We will also have a unit vector v_0 that is intended to represent the value 1. For a subset S of variables and a partial assignment $h : S \to \{0, 1\}$, we write the following SDP for the simultaneous MAX-2-SAT optimization problem:

$$\begin{array}{lrcll} & & \text{maximize} & t & \\ \text{s.t.} & \displaystyle\sum_{\substack{C \in \mathcal{C} \\ C = x_i \vee x_j}} \mathcal{W}_\ell(C) \cdot \left(\|v_0\|^2 - \frac{1}{4} \langle v_i - v_0, v_j - v_0 \rangle \right) & \geqslant & t & \forall \ell \in [k] \\ & \langle v_i - v_0, v_j - v_0 \rangle & \geqslant & 0 & \forall \text{ constraints } x_i \vee x_j \\ & \|v_i\|^2 & = & 1 & \text{for } i = -n, \ldots, n \\ & v_i & = & -v_{-i} & \text{for } i = 1, \ldots, n \\ & v_i & = & v_0 & \forall i \in S \text{ s.t. } h(i) = 1 \\ & v_j & = & -v_0 & \forall j \in S \text{ s.t. } h(j) = 0 \end{array}$$

Figure 14: Semidefinite Program (SDP) with a partial assignment $h : S \to \{0, 1\}$ for Simultaneous MAX-2-SAT

We first observe that for an optimal partial assignment h, the optimum of the above SDP is at least the optimum of the simultaneous maximization problem, by picking the solution $v_i = v_0$ if $x_i =$ TRUE, and $v_i = -v_0$ otherwise. For this vector solution, we have $1/4 \cdot \left(\|v_0\|^2 - \langle v_i - v_0, v_j - v_0 \rangle \right) = 1$ if the constraint $x_1 \vee x_2$ is satisfied by the assignment, and 0 otherwise. Since $\sum_{C \in \mathcal{C}} \mathcal{W}_\ell(C) = 1$ for all ℓ, the optimum of the SDP lies between 0 and 1.

Note that the rounding algorithm defined in [CMM06] does not depend on the structure of the vectors in the SDP solution. Thus, the following theorem that was proved without a partial assignment in [CMM06] also applies to above SDP.

Theorem C.3 *Given a single* MAX-2-SAT *instance ($k = 1$), there is an efficient randomized rounding algorithm such that, if the optimum of the above SDP is $1 - \varepsilon$, for n large enough, it returns an assignment such that the weight of the constraints satisfied is at least $1 - O(\sqrt{\varepsilon})$ in expectation.*

Now, using Markov's inequality, we can prove the following corollary.

Corollary C.4 *For k simultaneous instances of* MAX-2-SAT*, there is an efficient randomized rounding algorithm such that if the optimum of the above SDP is $1 - \varepsilon$, for n large enough, it returns an assignment that simultaneously satisfies at least $1 - O(k\sqrt{\varepsilon})$ fraction of the constraints in each instance with probability 0.9.*

Proof: We use the rounding algorithm given by Theorem C.3 to round a solution to the SDP for the k simultaneous instances that achieves an objective value of $1 - \varepsilon$. Observe that this solution is also a solution for the SDP for each of the instances by itself with the same objective value. Thus,

by Theorem C.3, for each of the instances, we are guaranteed to find an assignment such that the weight of the constraints satisfied is at least $1 - c_0\varepsilon$ in expectation, for some constant $c > 0$. Since, for any instance, the maximum weight an assignment can satisfy is at most 1, with probability at least $1 - 1/10 \cdot k$ for each instance, we get an assignment such that the weight of the constraints satisfied is at least $1 - 10ck \cdot \sqrt{\varepsilon}$. Thus, applying a union bound, with probability at least $1 - 1/10$, we obtain an assignment such that the weight of the satisfied constraints in *all* the k instances is at least $1 - 10ck \cdot \sqrt{\varepsilon}$. ■

Combining the above corollary with the reduction from any MAX-2-CSP to MAX-2-SAT, and the completeness of the SDP, we get a proof of Lemma B.1.

D Concentration inequalities

Lemma D.1 (McDiarmid's Inequality) *Let $X_1, X_2, \cdots, X_m$ be independent random variables, with X_i taking values in a set A_i for each i. Let $\mathsf{score} : \prod A_i \to \mathbb{R}$ be a function which satisfies:*

$$|\mathsf{score}(x) - \mathsf{score}(x')| \leqslant \alpha_i$$

whenever the vector x and x' differ only in the i-th co-ordinate. Then for any $t > 0$

$$\Pr[|\mathsf{score}(X_1, X_2, \cdots, X_m) - \mathbf{E}[\mathsf{score}(X_1, X_2, \cdots, X_m)]| \geqslant t] \leqslant 2\exp\left(\frac{-2t^2}{\sum_i \alpha_i^2}\right)$$

E The need for perturbing Opt

We construct 2 simultaneous instances of Max-1-SAT. Suppose the algorithm will picks at most r influential variables. Construct the two instances on $r + 1$ variables, with the weights of the variables decreasing geometrically, say, with ratio 1/3. The first instance requires all of them to be True, where as the second instance requires all of them to be False. Under a reasonable definition of "influential variables", the only variable left behind should the vertex with the least weight. We consider the Pareto optimal solution that assigns True to all but the last variable. If we pick the optimal assignment for the influential variables, and then randomly assign the rest of the variables, with probability 1/2, we get zero on the second instance.